



English |

config: cordis.yml apply a JSDoc cordis.yml cordis-
surface [system](#)
scripts/gen-config-catalog.ts pnpm run verify-config-catalog doc-sync ts config-catalog doc-typecheck
schemastery schema schema
Requires: inject cordis.yml packages/ [hmr](#) [policy](#)

@deepseek-ai/dsh-acp

agents · llm · sessionPersistence · sessions

```
/** Plugin config: the provider/model selection used for each ACP-created agent. */
export interface AcpConfig {
  /** Provider route for created agents. */
  provider?: string
  /** Model name for created agents. */
  model?: string
  /** Maximum summaries returned by one session/list page. */
  sessionListPageSize?: number
  /** Runtime-only transport override; production uses stdio. */
  stream?: Stream
}
```

Stream @agentclientprotocol/sdk

[packages/acp/acp/src/index.ts:75](#)

@deepseek-ai/dsh-agent-default-model

```
/** Composition entry for the default model selection. */
export interface Config {
  /** Registered provider route. */
  provider: string
  /** Provider-owned model id. */
  model: string
}
```

[packages/core/agent-default-model/src/index.ts:41](#)

@deepseek-ai/dsh-agent-instructions

sessionProjections

```
/** User-facing workspace instruction loader configuration. */
export interface Config {
  /** Harness home containing the fixed user-global `AGENTS.md`; defaults to `${DSH_HOME}` or `~/dsh`. */
  dshHome?: string
  /** Directory entries that identify the project root while walking upward from the session cwd. */
  projectRootMarkers?: string[]
  /** UTF-8 byte cap for one rendered baseline or dynamic batch; non-positive or non-finite disables loading. */
  maxBytes: number
  /** Maximum UTF-8 bytes read from one instruction file; larger files are ignored. */
  maxSourceBytes?: number
  /**
   * Ordered same-directory project candidates; every existing file loads, with
   * per-directory trimmed-content duplicates collapsed to the earliest candidate.
   */
  instructionFileCandidates?: string[]
  /**
   * Ordered same-directory local-overlay candidates loaded after the base files
   * under the same per-directory trimmed-content dedup; empty disables the overlay.
   */
  localInstructionFileCandidates?: string[]
}
```

[packages/context/agent-instructions/src/config.ts:18](#)

@deepseek-ai/dsh-agent-loop

agents · sessions · llm · tools · systemPrompt · sessionProjections

```
/** Agent-loop plugin configuration. */
export interface Config {
  /**
   * Maximum parallel-safe calls in flight per agent step. `1` is serial;
   * omission defaults to {@link DEFAULT_MAX_PARALLEL_TOOL_CALLS}.
   */
  maxParallelToolCalls?: number
  /** Agents created or resumed at plugin startup. */
  agents: (AgentOptions & {
    /** Stable config label used in logs and as the fresh combined-id prefix. */
    id: string
    /** Optional stable identity; remounts resume its materialized history, while first use creates it fresh. */
    sessionId?: SessionId
    /** Optional workspace for a fresh session. */
    cwd?: string
    /** Persisted session to resume instead of creating a fresh session. */
    resumeSessionId?: SessionId
  })[]
}
```

[AgentOptions](#) · [SessionId](#)

[packages/core/agent-loop/src/index.ts:311](#)

@deepseek-ai/dsh-agent-presets

loader · sessionProjections

```
/** Plugin config: which preset is the default, and where presets live. */
export interface Config {
  /** Preset id mounted when a caller names none. Missing at mount time fails loud. */
  default: string
  /** Scanned roots in precedence order; an earlier root wins a duplicate id. */
}
```

```

roots: PresetRoot[]

/**
 * Prepend this package's bundled shipped presets as a `system` root, before
 * every configured root, so the shipped set always mounts and wins a
 * duplicate id. The default survives a whole-`config` patch replacement;
 * only an explicit `false` — a deployment supplying purely its own presets,
 * or an embedder using the roster as bare machinery — drops the set.
 */
includeShippedRoot: boolean

/**
 * Append the harness home's `USER_PRESET_DIR` as a `user` root, after every
 * configured root. False mounts a roster without the derived writable root.
 */
includeUserRoot: boolean
}

/** One directory scanned for preset subdirectories. */
export interface PresetRoot {
  /** Directory holding one subdirectory per preset; a leading `~` expands. */
  path: string
  /** Trust recorded on every preset discovered under this root. */
  trust: PresetTrust
}

/**
 * Where a preset's composition came from. A `system` preset ships with the
 * deployment; a `user` preset was authored locally, by a person or by an
 * agent, and therefore carries the same trust as shell access.
 */
export type PresetTrust = 'system' | 'user'

```

[packages/preset/agent-presets/src/preset.ts:52](https://github.com/deepseek-ai/dsh-agent-tool/packages/preset/agent-presets/src/preset.ts:52)

@deepseek-ai/dsh-agent-tool- presentation

tools

```

/** Plugin config. */
export interface Config {
  /**
   * The form this agent's model sees. `native` sends every visible schema,
   * `ptc` sends only `run_code` plus a generated SDK, `both` sends both.
   * Required rather than defaulted: the deployment default is what a preset
   * without this row already gets, so an omitted value would mean the row was
   * composed for nothing.
   */
  mode: ToolPresentationMode
}

```

[ToolPresentationMode](#)

[packages/core/agent-tool-presentation/src/index.ts:38](#)

@deepseek-ai/dsh-api-gateway

typert

```

/** Gateway transport configuration. */
export interface Config {
  /** WebSocket Ping interval from 1 through 2,147,483,647 milliseconds. @default 2000 */
  readonly websocketHeartbeatIntervalMs?: number
}

```

[packages/api/gateway/src/index.ts:119](#)

@deepseek-ai/dsh-api-session-controller

agentDefaultModel · agents · attachments · llm · sessions · sessionProjections · sessionQuery · typert · workspaceRegistry

```

/** Session Controller deployment policy. */
export interface Config {
  /** Maximum cold Session artifact size eligible for one full projection observation. */
  readonly coldBlankProbeMaxBytes?: number
}

```

```
/** Override platform desktop-opener detection. */  
readonly nativeOpen?: boolean  
}
```

[packages/api/session-controller/src/index.ts:68](#)

@deepseek-ai/dsh-api-settings-controller

```
/** Native document-opening policy. */  
export interface Config {  
  /** Override platform desktop-opener detection. */  
  readonly nativeOpen?: boolean  
}
```

[packages/api/settings-controller/src/index.ts:36](#)

@deepseek-ai/dsh-attachment-local

```
/** Local attachment backend configuration. */  
export interface Config {  
  /** Explicit harness home; omitted follows `DSH_HOME`, then `~/.dsh`. */  
  dshHome?: string  
  /** Maximum encoded bytes accepted for one submitted image. Default: 20 MiB. */  
  maxImageBytes?: number  
  /** Maximum image count accepted in one submitted message. Default: 20. */  
  maxImagesPerMessage?: number  
  /** Maximum aggregate encoded image bytes accepted in one submitted message. Default: 200 MiB. */  
  maxMessageImageBytes?: number  
  /** Maximum intrinsic width multiplied by height accepted for one submitted image. Default: 64,000,000. */  
  maxImagePixels?: number  
  /** Maximum intrinsic width and maximum intrinsic height accepted for one submitted image. Default: 8192px. */  
  maxImageDimension?: number  
  /** Total-pixel budget of the stored provider-independent normalized image. */  
  normalizedImageMaxPixels?: number  
  /** Long-edge pixel cap of the stored provider-independent normalized image, applied after the total-pixel
```

```

budget. */
  normalizedImageMaxDimension?: number
/**
 * Encoded-byte target of the stored provider-independent normalized image;
 * the smallest quality-ladder output is kept when no quality fits.
 */
  normalizedImageMaxBytes?: number
/** Maximum simultaneous normalization or request-image transformations in this service instance. */
  imageCompressionConcurrency?: number
}

```

[packages/attachment/attachment-local/src/index.ts:55](#)

@deepseek-ai/dsh-bash-local

subprocess

```

/** Plugin config (all optional — `static Config` supplies the defaults). */
export interface Config {
  /** Default working directory for commands (default: process.cwd()). */
  cwd?: string
  /** Default foreground timeout in milliseconds. */
  timeoutMs?: number
  /** Upper bound for per-call timeout overrides. */
  maxTimeoutMs?: number
  /** Per-stream in-memory output cap; overflow spills to a temp file. */
  maxOutputBytes?: number
  /** Per-stream spill-file cap; larger streams retain only their in-memory tail. */
  maxSpillBytes?: number
  /** Grace period for kill escalation and inherited pipes; at most `MAX_TIMER_DELAY_MS`. */
  graceMs?: number
}

```

[packages/shell/bash-local/src/index.ts:41](#)

@deepseek-ai/dsh-bash-sandbox

subprocess · sandbox · sandboxPolicy

```
/**
 * Plugin config: the local executor's knobs, verbatim. The sandbox policy —
 * the default mode and fallback `workspace-write` root — is NOT here: it lives
 * on `ctx.sandboxPolicy` (`@deepseek-ai/dsh-sandbox-policy`), which resolves
 * each calling session's mode and cwd for every enforcing capability. The runner
 * choice is likewise the `ctx.sandbox` provider's config, not this executor's.
 */
export type Config = LocalConfig
```

[LocalConfig](#)

[packages/shell/bash-sandbox/src/index.ts:35](#)

@deepseek-ai/dsh-client-connection

webServer · credentials

```
/** Plugin config: the deployment's non-loopback serving authorities. */
export interface ConnectionConfig {
  /**
   * Authorities this deployment serves beyond loopback: exact `host:port`, or
   * port-less `host` matching any port. The /api trust fence refuses any
   * request whose Host is neither loopback nor listed here, so a
   * non-loopback (`0.0.0.0`) deployment must declare the names it is reached
   * by; the Web runtime derives LAN IP literals from an active all-interface
   * bind. An entry that is not a bare, canonical authority fails plugin load.
   */
  trustedHosts?: string[]
  /** Absolute browser-session lifetime in days. Default: 30. */
  cookieMaxAgeDays?: number
  /** Maximum buffered JSON body for every `/api` request. Default: 300 MiB. */
  maxRequestBodyBytes?: number
}
```

[packages/client/connection/src/index.ts:55](#)

@deepseek-ai/dsh-client-hmr

clientModules · webServer

```
/** Plugin config, validated by the same-named schemastery schema. */
export interface Config {
  /** Bundle stat-poll interval in milliseconds (default 500, the build-side watcher's polling default). */
  pollIntervalMs?: number
}
```

[packages/client/hmr/src/index.ts:31](#)

@deepseek-ai/dsh-code-runtime-worker-thread

```
/** Plugin config: every execution cap, changeable from `cordis.yml` (no hardcoded tunables). */
export interface Config {
  /**
   * Busy-time budget in milliseconds: the run fails with kind `timeout`
   * once the worker's MEASURED event-loop active time
   * (`worker.performance.eventLoopUtilization()`) exceeds this. Metering
   * measured busy time — not wall time, not host-side pending-call
   * bookkeeping — is what makes the budget both fair (a program awaiting a
   * slow tool accrues nothing) and ungameable (a hot loop accrues whether
   * or not a decoy dispatch is in flight).
   */
  computeMs?: number
  /**
   * Wall-clock ceiling in milliseconds; never pauses for anything. The
   * backstop for what busy-time cannot see (a program awaiting a promise
   * nobody will resolve). At most `2_147_483_647` (Node's maximum
   * `setTimeout` delay, about 24.9 days): a longer value is rejected at load
   * because `setTimeout` would clamp it to 1 ms.
   */
  maxWallMs?: number
}
```

```
* Hard cap for serialized log-array, completion-value, and failure-message payloads;
* fixed result-envelope syntax is excluded.
*/
maxOutputBytes?: number
/** The worker's max old-generation heap in MiB (`resourceLimits`); overflow kills the worker, surfacing as kind
`worker-exit`. */
maxOldGenerationSizeMb?: number
}
```

[packages/code-runtime/code-runtime-worker-thread/src/index.ts:25](#)

@deepseek-ai/dsh-compaction-basic

llm · tokenMeter · sessions

```
/** Basic compaction configuration with an optional exact-target policy table. */
export interface BasicCompactionConfig extends CompactionPolicyConfig {
  /** Exact provider/model overrides; duplicate targets fail plugin load. */
  modelPolicies?: ModelCompactPolicyConfig[]
  /** Enable automatic step-boundary pressure and overflow-recovery listeners. Defaults to `true`. */
  auto?: boolean
}

/** Policy fields shared by the default policy and exact model overrides. */
export interface CompactionPolicyConfig {
  /** Compact at this fraction of the model's context window. Defaults to `0.8`. */
  thresholdRatio?: number
  /** Recent context retained as a fraction of the model's window. Defaults to `0.16`. */
  retainRatio?: number
  /** Absolute recent-context budget; mutually exclusive with `retainRatio`. */
  retainTokens?: number
  /** Summary provider; set together with `summarizationModel`, or inherit the conversation target. */
  summarizationProvider?: string
  /** Summary model; set together with `summarizationProvider`, or inherit the conversation target. */
  summarizationModel?: string
  /** Provider generation cap for summarization. Defaults to `8192`. */
  maxTokens?: number
  /** Extra attempts after the first compaction when pressure remains above threshold. Defaults to `1`. */
}
```

```

compactionRetries?: number

/** Maximum retries after canonical context overflow; `0` disables recovery. Defaults to `1`. */
maxOverflowRetries?: number
}

/** Exact provider/model override merged over the default compaction policy. */
export interface ModelCompactPolicyConfig extends CompactionPolicyConfig {
  /** Registered provider route to match. */
  provider: string
  /** Exact routed model id to match within `provider`. */
  model: string
}

```

[packages/compaction/compaction-basic/src/types.ts:38](#)

@deepseek-ai/dsh-compaction-tool-result-pruner

tokenMeter

```

/** Character-budget policy for deterministic tool-result pruning. */
export interface ToolResultPruneConfig {
  /** Prune when total text exceeds this many Unicode code points. Defaults to `8192`. */
  thresholdChars?: number
  /** Maximum leading Unicode code points retained. Defaults to `4096`. */
  headChars?: number
  /** Maximum trailing Unicode code points retained. Defaults to `1024`. */
  tailChars?: number
}

```

[packages/compaction/compaction-tool-result-pruner/src/types.ts:5](#)

@deepseek-ai/dsh-cordis-host-runner

tools

```
/** Runner configuration. */  
export interface Config {  
  /** Maximum synchronous VM evaluation time in milliseconds. */  
  vmTimeoutMs?: number  
}
```

[packages/extensions/cordis-host-runner/src/index.ts:88](#)

@deepseek-ai/dsh-credentials-local

```
/** Plugin config: file location and hot-reload behavior. */  
export interface Config {  
  /** Credentials document path; defaults to `.credentials.yaml` under the harness home. */  
  path?: string  
  /** Harness home used when `path` is omitted; defaults to `$DSH_HOME` or `~/dsh`. */  
  dshHome?: string  
  /** Watch the document and hot-publish external edits; defaults to true. */  
  watch?: boolean  
  /** Watcher write-settle window in milliseconds; defaults to 100. */  
  debounceMs?: number  
}
```

[packages/credentials/credentials-local/src/index.ts:64](#)

@deepseek-ai/dsh-e2b

```
/** Configuration for the shared E2B sandbox owner. */  
export interface Config {  
  /** API key; omission reads `E2B_API_KEY`. It is never forwarded into the sandbox. */  
  apiKey?: string  
  /** Shared remote working directory, created before adapters receive the sandbox. */  
  cwd?: string  
  /** E2B sandbox lifetime in milliseconds; expiry always deletes the sandbox. */  
  timeoutMs?: number  
}
```

[packages/e2b/e2b/src/index.ts:43](#)

@deepseek-ai/dsh-experimental-agent-team

[agents](#) · [sessions](#) · [sessionPersistence](#) · [sessionProjections](#) · [subagents](#)

```
/** Team-service deployment limits. */
export interface Config {
  /** Maximum immutable teammate names retained by one Team. */
  readonly maxMembers?: number
  /** Maximum non-deleted tasks retained by one Team. */
  readonly maxTasks?: number
  /** Maximum queued-minus-delivered messages for one target member. */
  readonly maxPendingMessagesPerMember?: number
  /** Maximum UTF-8 bytes in one complete sender-framed delivery. */
  readonly maxMessageBytes?: number
  /** Maximum milliseconds allowed for Team-owned runtime disposal. */
  readonly disposalTimeoutMs?: number
}
```

[packages/experimental/agent-team/src/types.ts:125](#)

@deepseek-ai/dsh-experimental-code-runtime-python

```
/** Plugin config: every cap, changeable from `cordis.yml` (no hardcoded tunables). */
export interface Config {
  /**
   * RLIMIT_CPU in whole seconds (a positive integer — `setrlimit` in the child
   * rejects a float). The child sets the soft limit to `cpuSeconds` and the
   * hard limit to `cpuSeconds + 1`: the kernel delivers SIGXCPU at the soft
   * limit, which the host classifies as a `timeout`; the +1s hard limit is a
   * SIGKILL backstop for a program that traps SIGXCPU. Granularity is seconds —
```

```

* a coarser counterpart to the worker backend's millisecond `computeMs`.
*/
cpuSeconds?: number
/** Wall-clock ceiling in milliseconds; backstops CPU time for programs awaiting a promise nobody resolves. */
maxWallMs?: number
/**
* RLIMIT_AS in mebibytes; caps address space so a runaway allocation fails
* cleanly. Not applied on Darwin, where the dyld shared cache mapped into
* every process at exec exceeds any practical cap and the kernel rejects
* the call; `cpuSeconds` and `maxWallMs` still bound the run there. Bounds
* `maxLogBytes`/`maxValueBytes` at load on EVERY platform (this static check
* runs on Darwin too, where only the runtime `setrlimit` is skipped): each
* budget times a worst-case Unicode expansion must fit this byte count minus a
* fixed interpreter baseline, so a near-budget output cannot breach the address
* space during the child's build-and-encode.
*/
addressSpaceMb?: number
/**
* Shared byte budget for captured log text (host-side ledger). Bounded at load
* against `addressSpaceMb`: the child builds and encodes a near-budget entry
* under RLIMIT_AS with several copies live at once, so this cap times the
* worst-case Unicode expansion must fit the address space left after the
* interpreter baseline (see `addressSpaceMb`) — a load-time rejection, not a
* runtime clamp. Also bounded at load by the host's configured heap like
* `maxValueBytes` (see its JSDoc): the effective frame cap minus the frame
* envelope.
*/
maxLogBytes?: number
/**
* Byte cap for the completion value. Bounded at load against `addressSpaceMb`
* the same way `maxLogBytes` is: the child builds and encodes a near-budget
* value under RLIMIT_AS with several copies live at once, so this cap times the
* worst-case Unicode expansion must fit the address space left after the
* interpreter baseline. Both budgets are ALSO bounded at load by the host's
* configured heap: the effective frame cap (the protocol cap, or a lower
* heap-derived ceiling when the host heap cannot safely parse a near-cap
* frame — see `hostFrameParseCeiling`) minus the frame envelope, so a budget
* whose honest frame could OOM the host's own JSON.parse is rejected up
* front.
*/

```

```

maxValueBytes?: number

/** SIGTERM→SIGKILL grace period on kill, matching bash-local's default. */
graceMs?: number

/**
 * Absolute path, relative path, or basename of a CPython 3.10+ interpreter.
 * Resolved and validated once at plugin load under a five-second force-kill
 * deadline; a basename searches `PATH`.
 */
pythonBin?: string
}

```

[packages/experimental/code-runtime-python/src/index.ts:42](https://github.com/deepseek-ai/dsh-experimental-inspector/blob/main/packages/experimental/code-runtime-python/src/index.ts#L42)

@deepseek-ai/dsh-experimental-inspector

webServer

```

/** Host plugin configuration. Fetch capture is enabled by default. */
export interface Config extends Omit<InspectorOptions, 'clientOrigins'> {
  /** Browser origins allowed to open the Client ingest WebSocket. */
  clientOrigins?: string[]
}

/** User-facing Host options; every memory and lifecycle bound is configurable. */
export interface InspectorOptions {
  /** Loopback address used by the Worker HTTP and WebSocket endpoint. */
  readonly host?: '127.0.0.1'

  /** First port to bind; occupied ports advance until one is available. */
  readonly port?: number

  /** Additional exact browser origins admitted to the Client ingest socket. */
  readonly clientOrigins?: readonly string[]

  /** Whether to observe calls made through the current global fetch function. */
  readonly captureFetch?: boolean

  /** Maximum request-body prefix retained for one fetch. */
  readonly maxRequestBodyBytes?: number

  /** Maximum response-body prefix retained for one fetch. */
}

```

```

readonly maxResponseBodyBytes?: number
/** Maximum raw bytes encoded into one body observation. */
readonly maxBodyChunkBytes?: number
/** Maximum total request and response body bytes retained by the Worker. */
readonly maxJournalBytes?: number
/** Maximum active and completed fetch requests retained by the Worker. */
readonly maxRetainedRequests?: number
/** Maximum encoded bytes accepted in one source transport frame. */
readonly maxSourceFrameBytes?: number
/** Maximum observation records accepted in one source batch. */
readonly maxSourceRecordsPerFrame?: number
/** Maximum records waiting in one producer queue. */
readonly maxQueuedRecords?: number
/** Maximum encoded bytes waiting in one producer queue. */
readonly maxQueuedBytes?: number
/** Maximum time allowed for the Worker to become ready. */
readonly startupTimeoutMs?: number
/** Grace period before a stopping Worker is terminated. */
readonly stopTimeoutMs?: number
/** Initial upper bound for randomized Client reconnect delay. */
readonly clientReconnectBaseMs?: number
/** Maximum upper bound for randomized Client reconnect delay. */
readonly clientReconnectMaxMs?: number
/** Deadline for one Worker-to-Client Runtime or Sources request. */
readonly clientRuntimeTimeoutMs?: number
/** Deadline for one non-CDP semantic query. */
readonly queryTimeoutMs?: number
/** Maximum live object handles retained per Client Runtime session. */
readonly maxClientRuntimeObjects?: number
/** Maximum descriptors returned by one Client property request. */
readonly maxClientRuntimeProperties?: number
/** Maximum encoded bytes read for one Client script or source map. */
readonly maxClientSourceBytes?: number
/** Maximum Context and Fiber nodes retained in one realm snapshot. */
readonly maxCordisNodes?: number
/** Disconnected Cordis snapshots retained after their live realm closes. */
readonly maxDisconnectedCordisTrees?: number
}

```


@deepseek-ai/dsh-experimental-tool-agent-team

agents · agentTeams · tools · systemPrompt

```
/** Tool routing configuration. */  
export interface Config {  
  /** Continuable-subagent provider used for fresh teammates. */  
  readonly freshProvider?: string  
  /** Continuable-subagent provider used for completed-prefix fork teammates. */  
  readonly forkProvider?: string  
}
```

[packages/experimental/tool-agent-team/src/index.ts:17](#)

@deepseek-ai/dsh-file-reference-local

agents · sessionProjections

```
/** Local file-reference discovery configuration. */  
export interface Config {  
  /** Maximum ranked candidates returned for one query. */  
  maxResults?: number  
  /** Maximum indexed files and directories per agent workspace. */  
  maxEntries?: number  
  /** Directory basenames never traversed or offered. */  
  excludedDirectories?: string[]  
}
```

[packages/context/file-reference-local/src/index.ts:34](#)

@deepseek-ai/dsh-fs-local

```
/** Configuration for the local filesystem backend. */
export interface Config {
  /** Base directory for relative paths. Defaults to `process.cwd()`. */
  cwd?: string
  /**
   * Exclusive UTF-8 byte limit on each overwrite-diff side, capped by the
   * runtime's safe allocation/decode maximum. Defaults to 10 MiB.
   */
  diffBasisMaxBytes?: number
}
```

[packages/fs/fs-local/src/index.ts:41](#)

@deepseek-ai/dsh-fs-sandbox

sandboxPolicy

```
/**
 * Plugin config: the local backend's knobs verbatim (`cwd` resolution default
 * and `diffBasisMaxBytes` overwrite-presentation bound). The sandbox default
 * (mode + `workspace-write` fallback root) is NOT here — `ctx.sandboxPolicy`
 * resolves each calling session for every enforcing capability.
 */
export type Config = LocalConfig
```

[LocalConfig](#)

[packages/fs/fs-sandbox/src/index.ts:45](#)

@deepseek-ai/dsh-goal

agents · sessionProjections

```
/** Deployment defaults for goal creation. */
export interface Config {
  /** Total rounds used when a create request omits its own cap. */
  defaultMaxGoalRounds?: number
}
```

```
}
```

[packages/goal/goal/src/index.ts:172](#)

@deepseek-ai/dsh-headless

agentDefaultModel · agents · sessions

```
/** Plugin config: the task resolved from this app's injected provider service. */
export interface Config {
  /** The prompt text for the single run. */
  task: string
}
```

[packages/bundle/headless/src/index.ts:33](#)

@deepseek-ai/dsh-hooks-claude-code

shell · sessionProjections

```
/** Plugin config: where the CC hook config lives + substitution roots. */
export interface Config {
  /**
   * Path to a `hooks.json` or a settings file whose `hooks` key holds the config.
   * Process-level: read once at load, a relative path resolves against the process
   * launch cwd, so one config applies to the whole process.
   * TODO(per-session-hook-config): per-session discovery of a project-local
   * `hooks.json` from each `session/new.cwd`.
   */
  configPath: string
  /**
   * Replaces `${CLAUDE_PLUGIN_ROOT}` in command strings (the plugin's root dir).
   */
  pluginRoot?: string
  /**
   * Replaces `${CLAUDE_PROJECT_DIR}` in command strings AND is exported as the
   * `CLAUDE_PROJECT_DIR` env var for hook processes. When omitted, the env var
```

```

* defaults per-run to the agent's session workspace (`session.header.cwd`, the
* same dir the hook runs in) — Claude Code always exports this var, and common
* unmodified hooks reference `$CLAUDE_PROJECT_DIR` for project-relative paths.
*/
projectDir?: string
/** Default per-hook timeout in ms when a hook sets none (CC default: 600000). */
defaultTimeoutMs?: number
/** Character cap for the `hook/result` event's persisted stderr summary. */
stderrSummaryMaxChars?: number
}

```

[packages/hooks/hooks-claude-code/src/index.ts:46](https://github.com/deepseek-ai/dsh-hooks-codex/blob/main/packages/hooks/hooks-claude-code/src/index.ts#L46)

@deepseek-ai/dsh-hooks-codex

shell · sessionProjections

```

/** Plugin config: where the Codex hooks.json lives + the model name for payloads. */
export interface Config {
  /**
   * Path to a Codex `hooks.json`. Process-level: read once at load, a relative
   * path resolves against the process launch cwd.
   * TODO(per-session-hook-config): per-session project-local discovery from each
   * `session/new.cwd`.
   */
  configPath: string
  /** The model name stamped on every payload (Codex includes `model` on each event). */
  model?: string
  /** Default per-hook timeout in ms when a hook sets none (Codex default: 600000). */
  defaultTimeoutMs?: number
  /** Character cap for the `hook/result` event's persisted stderr summary. */
  stderrSummaryMaxChars?: number
}

```

[packages/hooks/hooks-codex/src/index.ts:45](https://github.com/deepseek-ai/dsh-hooks-codex/blob/main/packages/hooks/hooks-codex/src/index.ts#L45)

@deepseek-ai/dsh-host-directory-picker-browse

```
/** Validated plugin configuration. */  
export interface Config {  
  /** Complete-result bound of one listing level; see {@link BrowseDirectoryPicker.Config}. */  
  maxEntries: number  
}
```

[packages/host/directory-picker-browse/src/index.ts:181](#)

@deepseek-ai/dsh-host-frontend-static

webServer · connection

```
/** Plugin config: the dist anchor. */  
export interface Config {  
  /** Absolute path of index.html inside the dist root. */  
  distIndex: string  
}
```

[packages/host/frontend-static/src/index.ts:30](#)

@deepseek-ai/dsh-host-webserver

```
/** Web server listen and response-compression config. */  
export interface Config {  
  /** Listen host; the two supported values are loopback and all-interfaces. */  
  host: '127.0.0.1' | '0.0.0.0'  
  /** Listen port; zero requests an OS-assigned port. */  
  port: number  
  /** Response compression for socket-backed HTTP requests. @default 'none' */  
  compression?: 'none' | 'gzip'  
  /** Gzip DEFLATE level from 0 through 9. @default 1 */
```

```
compressionLevel?: number

/** Minimum known response length eligible for gzip; unknown-length streams are eligible. @default 1024 */
compressionThresholdBytes?: number
}
```

[packages/host/webserver/src/index.ts:59](#)

@deepseek-ai/dsh-invariants

```
/** Runtime invariant selection configured on the service plugin. */
export interface Config {
  /** Global switch; defaults to `true`. */
  readonly enabled?: boolean

  /** Case-sensitive JavaScript regex sources that admit package names; empty admits all. */
  readonly package_allowlist?: string[]

  /** Case-sensitive JavaScript regex sources that exclude package names after allowlist matching. */
  readonly package_blocklist?: string[]
}
```

[packages/runtime-diagnostics/invariants/src/index.ts:15](#)

@deepseek-ai/dsh-jobs-local

```
/** Configuration for the process-local job registry. */
export interface Config {
  /**
   * Maximum `running` plus `stopping` jobs per exact owner or in the shared unowned bucket;
   * omission defaults to 10.
   */
  readonly maxConcurrentJobsPerOwner?: number
}
```

[packages/jobs/jobs-local/src/index.ts:31](#)

@deepseek-ai/dsh-llm-deepseek

llm

```
/**
 * Plugin config, validated by the same-named schemastory schema and doubling
 * as the `llm-deepseek` settings-section shape. Every field is optional in
 * yaml: a missing API key resolves through {@link Config.apiKeyEnv} at each
 * request (a request without any key fails with `MISSING_CREDENTIAL`, not at
 * plugin load), omitted thinking mode uses the provider default, and omitted
 * reasoning effort resolves to `high`.
 */
export interface Config {
  /** Credential reference (environment-variable name) resolved per request; defaults to `DEEPSEEK_API_KEY`. */
  apiKeyEnv?: string
  /** Endpoint base; falls back to $DEEPSEEK_BASE_URL from a trusted environment layer, then the public API. */
  baseUrl?: string
  /** Deployment thinking policy; `disabled` limits every conversation request to `off`. */
  thinking?: 'enabled' | 'disabled'
  /** Default thinking effort (default `high`); `off` disables thinking per request. */
  reasoningEffort?: 'off' | 'low' | 'high' | 'max'
  /** Default per-request output cap (default 256,000); a model's own cap and explicit request values win. */
  maxTokens?: number
  /** Positive context capacity used when the selected model has no exact value (default 1,000,000). */
  defaultContextWindow?: number
  /** Advisory models shown by discovery consumers; defaults to V4 Flash, V4 Pro, and V4 Flash Vision Exp. */
  models?: DeepSeekCatalogModel[]
  /** Maximum provider idle time while one stream read is outstanding (default five minutes). */
  streamIdleTimeoutMs?: number
  /** Maximum accumulated file-referenced image bytes per chat request (default 128 MiB). */
  maxRequestFilesBytes?: number
  /** Maximum accumulated base64 image payload after Files API fallback (default 20 MiB). */
  maxInlineRequestImageBytes?: number
  /** Maximum number of represented images per chat request (default 600). */
  maxImagesPerRequest?: number
  /** Raw-byte removal step after the request exceeds its file bound (default 64 MiB). */
  imageOffloadByteQuantum?: number
  /** Base64-byte removal step after inline fallback exceeds its bound (default 10 MiB). */
  inlineImageOffloadByteQuantum?: number
}
```

```

/** Image-count removal step after the request exceeds its count bound (default 20). */
imageOffloadCountQuantum?: number

/** Maximum duration of one request-image Files API resolution (default one minute). */
filesApiTimeoutMs?: number

/** Explicit lifetime assigned to each uploaded image (default seven days). */
fileExpiresAfterSeconds?: number

/** Remaining lifetime below which an indexed file is replaced (default one hour). */
fileRefreshMarginSeconds?: number

/** Oldest harness-owned files deleted before one quota-recovery upload retry (default 100). */
fileQuotaCleanupBatch?: number

/** Provider-owned model-request retry policy; omission uses normal mode with five retries. */
retryPolicy?: RetryPolicyConfig
}

/** One optional model entry advertised by the direct-fetch adapter. */
export interface DeepSeekCatalogModel {
  /** Wire model id accepted by the configured endpoint. */
  id: string

  /** Selector label; defaults to {@link id}. */
  name?: string

  /** Optional selector detail for deployments with similar model variants. */
  description?: string

  /** Known combined request/response context capacity; omitted when deployment metadata is unavailable. */
  contextWindow?: number

  /** Per-request output cap for this model; omission falls back to the profile's {@link
DeepSeekConnectionOptions.maxTokens}. */
  maxTokens?: number

  /** Accepted request modalities; omission is text-only. */
  inputModalities?: ModelModality[]

  /** Total-pixel budget for one deterministic request preview, or the 512-by-512 `low` preset. */
  imagePixelBudget?: number | 'low'

  /** Encoded-byte target for one deterministic request preview; the smallest quality-ladder output is used when
no quality fits. */
  imageMaxBytes?: number
}

```

[ModelModality](#) · [RetryPolicyConfig](#)

[packages/llm/llm-deepseek/src/index.ts:125](#)

@deepseek-ai/dsh-llm-pi-ai

llm

```
/** Plugin configuration: the provider routes this instance owns. */
export interface Config {
  /**
   * pi-ai provider routes, keyed by provider. An empty (or omitted) dict is
   * the dormant settings-driven posture: the adapter mounts with no routes
   * and registers them the moment a settings section supplies profiles.
   */
  providers?: Record<string, PiAiProviderProfile>
}

/** Configuration for one pi-ai provider route; the `providers` dict key IS the route. */
export interface PiAiProviderProfile {
  /** Credential reference (environment-variable name) resolved per request through `ctx.credentials`. */
  apiKeyEnv?: string
  /** Name shown by configuration surfaces; defaults to the route key. */
  displayName?: string
  /**
   * Wire protocol every model on this route speaks. Omission keeps each
   * installed catalog model's own protocol, which is why a catalog route needs
   * no protocol at all; a route the catalog does not ship must name one.
   */
  api?: string
  /** Endpoint for this route's models; defaults to the installed catalog's endpoint. */
  baseUrl?: string
  /**
   * This route's model catalog. Omission serves the installed catalog for the
   * route unchanged; an explicit list replaces it, each entry defaulting its
   * unset fields from the installed model of the same id.
   */
  models?: PiAiModelProfile[]
  /**
   * Installed-catalog customizations by model id: each entry reshapes that
   * one model with the same fields a {@link models} entry takes, while the
   * rest of the catalog keeps serving untouched. Only meaningful on a catalog
```

```

* route with no `models` list — `models` already replaces the catalog, so
* an override beside it, on a route the catalog does not ship, or naming a
* model the catalog does not describe is refused rather than skipped.
*/
modelOverrides?: Record<string, PiAiModelOverride>
/**
* pi-ai wire-compatibility switches defaulting every model on this route
* whose protocol declares them; each model's own `compat` overrides per
* field. What neither sets keeps the installed catalog entry's value, then
* pi-ai's own detection. A switch no model on the route could read is
* refused rather than left looking applied.
*/
compat?: PiAiCompatProfile
/**
* Context capacity for a model this route lists that neither the entry nor
* the installed catalog sizes (default 262,144). A guess by construction, so
* a deployment whose gateway serves smaller models corrects it here.
*/
defaultContextWindow?: number
/**
* Output capability for a model this route lists that neither the entry nor
* the installed catalog sizes (default 32,768). This sizes the model; it
* never becomes a per-request cap on its own.
*/
defaultMaxTokens?: number
/**
* Request modalities for a model this route lists that neither its entry's
* {@link PiAiModelProfile.input} nor the installed catalog declares (default
* `[text]`). A fallback like the capacities above, not an override: a
* catalog model keeps the modalities the catalog records for it, and this
* value never narrows one. A gateway serving vision models the catalog does
* not describe declares `[text, image]` once here instead of on every entry.
* Unlike an entry's list, this one may not be empty — nothing sits below it
* to answer instead.
*/
defaultInput?: PiAiModality[]
/** Provider request headers, validated against Fetch when the profile resolves; Harness attribution wins
reserved names. */
headers?: Record<string, string>
/** Provider-neutral pi-ai reasoning level. */

```

```

reasoning?: ModelThinkingLevel
/** Token budgets used by reasoning providers that support them. */
thinkingBudgets?: ThinkingBudgets
/** Prompt-cache retention preference. */
cacheRetention?: CacheRetention
/** Streaming transport preference. */
transport?: Transport
/** HTTP/provider SDK timeout in milliseconds. */
timeoutMs?: number
/** WebSocket connection timeout in milliseconds. */
websocketConnectTimeoutMs?: number
/** Maximum provider idle time while one stream read is outstanding. */
streamIdleTimeoutMs?: number
/**
 * Maximum base64-encoded image payload per request. When a request's
 * accumulated images exceed it, the oldest images are replaced by text
 * placeholders until the request fits, so a long session keeps completing
 * requests instead of being rejected by a request-size cap.
 */
maxRequestImageBytes?: number
/** Total-pixel budget for each deterministic inline request version. */
requestImagePixelBudget?: number
/**
 * Raw encoded-byte target for each deterministic inline request version;
 * the smallest quality-ladder output is used when no quality fits.
 */
requestImageMaxBytes?: number
/** Provider-owned model-request retry policy; omission uses normal mode with five retries. */
retryPolicy?: RetryPolicyConfig
}

/** One configured model entry: an id plus the catalog fields it overrides. */
export interface PiAiModelProfile {
  /** Model id sent to the provider and accepted by {@link GenerateOptions.model}. */
  id: string
  /** Display name for selectors; defaults to the catalog name, then the id. */
  name?: string
  /** Maximum combined request and response context in tokens. */
  contextWindow?: number
}

```

```

* Maximum output tokens. Configuring one also makes it this model's
* per-request default; a value inherited from the installed catalog, or the
* route's fallback, is the model's capability and never becomes a request
* default on its own.
*/
maxTokens?: number

/**
* Request modalities this model accepts. Absent — or empty, which describes
* a model that accepts nothing and so states no answer either — keeps the
* installed catalog entry's modalities, then the route's `defaultInput`.
* Declaring images is what makes a hand-declared vision model usable, and
* declaring text alone corrects a catalog model whose gateway does not serve
* what the catalog records. This is a claim about the endpoint, not a check
* of it: nothing interrogates a gateway for what it accepts, so a model
* claiming images its endpoint refuses is refused by the provider instead,
* mid-turn.
*/
input?: PiAiModality[]

/**
* Selectable reasoning efforts. Absent inherits the installed catalog
* entry's capability (a hand-declared model has none and does not reason);
* `false` declares a non-reasoning model, which is how a profile strips
* reasoning from a catalog model its gateway cannot serve; a non-empty dict
* declares the offered levels and their wire spellings.
*/
reasoningEfforts?: false | PiAiReasoningEfforts

/** pi-ai wire-compatibility switches for this model, winning over the route's per field; one its protocol does not
declare is refused. */
compat?: PiAiCompatProfile
}

/**
* Customization of one installed catalog model, keyed by its id in the
* route's `modelOverrides` dict — the same fields a `models` entry may set,
* with the id living in the key. Unlike a `models` list, overrides leave the
* rest of the catalog serving untouched, which is what makes "correct one
* model, keep the other thirty-seven" a three-line edit.
*/
export type PiAiModelOverride = Omit<PiAiModelProfile, 'id'>

```

```

/**
 * pi-ai wire-compatibility switches, set on the route (its models' default) or
 * per model (winning over the route, field by field).
 *
 * pi-ai decides each of these from the provider id and baseURL when no layer
 * sets it, and a private gateway's URL says nothing: for an endpoint it does
 * not recognize the detection answers as though it were OpenAI itself, which
 * is wrong for most OpenAI-compatible gateways. So every field here is one a
 * deployment must be able to state because nothing can infer it, while the
 * fields pi-ai's catalog sets for a named vendor stay withheld.
 *
 * A field belongs to the protocols whose upstream compat type declares it: a
 * model-level switch its protocol does not take fails resolution, and a
 * route-level one skips past models it cannot fit. "The three Responses
 * protocols" below means `openai-responses`, `azure-openai-responses`, and
 * `openai-codex-responses`, which pi-ai gives one shared compat type, so a
 * switch settable on one is settable on all three.
 */
export interface PiAiCompatProfile {
  /** Whether the endpoint accepts `store`; `openai-completions`. */
  supportsStore?: boolean
  /**
   * Whether the endpoint accepts the `developer` role for the system prompt,
   * which pi-ai sends only to a reasoning model; `false` keeps `system`.
   * `openai-completions` and the three Responses protocols.
   */
  supportsDeveloperRole?: boolean
  /** Whether the endpoint accepts `reasoning_effort`; `openai-completions`. */
  supportsReasoningEffort?: boolean
  /** Whether the endpoint accepts `stream_options: {include_usage: true}`; `openai-completions`. */
  supportsUsageInStreaming?: boolean
  /**
   * Whether streams include `finish_reason`; `false` lets pi-ai infer the
   * terminal reason when the stream ends; `openai-completions`.
   */
  supportsFinishReason?: boolean
  /** Which output-cap field the endpoint reads; `openai-completions`. */
  maxTokensField?: NonNullable<OpenAICompletionsCompat['maxTokensField']>
  /** Whether tool results must carry `name`; `openai-completions`. */
  requiresToolResultName?: boolean

```

```
/** Whether a user message after tool results needs an assistant message between; `openai-completions`. */
requiresAssistantAfterToolResult?: boolean

/** Whether thinking blocks must travel as text in `<thinking>` delimiters; `openai-completions`. */
requiresThinkingAsText?: boolean

/** Whether replayed assistant messages need an empty `reasoning_content` while reasoning is on; `openai-completions`. */
requiresReasoningContentOnAssistantMessages?: boolean

/** Reasoning parameter format the endpoint expects; `openai-completions`. */
thinkingFormat?: PiAiThinkingFormat

/**
 * Kwargs sent as `chat_template_kwargs`, which pi-ai reads only under the
 * two `chat-template` thinking formats; `openai-completions`. Nothing checks
 * that pairing: the format in force may come from the installed catalog
 * entry or from pi-ai's own baseURL detection, neither of which resolution
 * can read, so kwargs set beside another format are sent nowhere.
 */
chatTemplateKwargs?: NonNullable<OpenAICompletionsCompat['chatTemplateKwargs']>

/** Arguments sent as `chat_template_args` under the `baseten` thinking format; `openai-completions`. */
chatTemplateArgs?: NonNullable<OpenAICompletionsCompat['chatTemplateArgs']>

/** Whether the endpoint accepts `thinking_token_budget` to cap vLLM reasoning; `openai-completions`. */
supportsThinkingTokenBudget?: boolean

/**
 * Whether the endpoint accepts `strict` in tool definitions;
 * `openai-completions`, the three Responses protocols, `bedrock-converse-stream`.
 */
supportsStrictMode?: boolean

/** Prompt-cache marker convention; `openai-completions`. */
cacheControlFormat?: NonNullable<OpenAICompletionsCompat['cacheControlFormat']>

/**
 * Whether the endpoint accepts long prompt-cache retention;
 * `openai-completions`, the three Responses protocols, `anthropic-messages`.
 */
supportsLongCacheRetention?: boolean

/** Whether the endpoint accepts per-tool `eager_input_streaming`; `anthropic-messages`. */
supportsEagerToolInputStreaming?: boolean

/** Whether the endpoint accepts `cache_control` on tool definitions; `anthropic-messages`. */
supportsCacheControlOnTools?: boolean

/** Whether the endpoint accepts the `temperature` request field; `anthropic-messages`. */
supportsTemperature?: boolean

/** Whether to force adaptive thinking regardless of model id; `anthropic-messages`. */
```

```

forceAdaptiveThinking?: boolean

/** Whether to replay an empty thinking signature instead of converting thinking to text; `anthropic-
messages`. */
allowEmptySignature?: boolean

/** Whether the endpoint accepts Anthropic strict tool schemas; `anthropic-messages`. */
supportsStrictTools?: boolean
}

/** One request modality a pi-ai model may accept. */
export type PiAiModality = Model<Api>['input'][number]

/**
 * Selectable reasoning efforts for one model: each key is a level the model
 * offers (and selectors show), and its value is the wire spelling dispatch
 * sends for it. `off` alone may leave its value empty — "supported, send
 * nothing" — because for most providers not thinking is the parameter's
 * absence; every other declared level must name a wire value. A level absent
 * from the dict is not offered.
 */
export type PiAiReasoningEfforts = Partial<Record<ModelThinkingLevel, string | null>>

/** One reasoning-dispatch wire format a profile may name. */
export type PiAiThinkingFormat = NonNullable<OpenAICompletionsCompat['thinkingFormat']>

```

```

Api @earendil-works/pi-ai < CacheRetention @earendil-works/pi-ai < Model @earendil-works/pi-ai <
ModelThinkingLevel @earendil-works/pi-ai < OpenAICompletionsCompat @earendil-works/pi-ai < RetryPolicyConfig
· ThinkingBudgets @earendil-works/pi-ai < Transport @earendil-works/pi-ai )

```

[packages/llm/llm-pi-ai/src/config.ts:213](https://github.com/earendil-works/pi-ai/blob/main/packages/llm/llm-pi-ai/src/config.ts#L213)

@deepseek-ai/dsh-llm-replay

llm

```

/** Plugin config: the {@link ReplayConfig} inputs, each defaulting to its `DSH_SNAPSHOT_*` env var in `apply`.
 */
export interface Config {
  /** Override the fixture path; defaults to `$DSH_SNAPSHOT_FILE`. */
  file?: string
}

```

```

/** Override the sidecar path; defaults to `$_DSH_SNAPSHOT_OVERRIDE`. */
overrideFile?: string
/**
 * Override the child-log paths; defaults to `$_DSH_SNAPSHOT_CHILD_FILES` (a
 * path-separator-delimited list). Each is a recorded subagent session log for
 * a nested-agent scenario; absent/empty for a single-session scenario.
 */
childFiles?: string[]
/** Optional replay-only provider catalog; absent or empty selects catch-all waterfall replay. */
providers?: ReplayProviderConfig[]
/** Optional per-chunk pacing delay in ms (see {@link ReplayConfig.paceMs}); absent keeps burst yield. */
paceMs?: number
}

/** One provider route exposed by the replay adapter. */
export interface ReplayProviderConfig {
  /** Provider route used for replay requests. */
  id: string
  /** Selector label; defaults to {@link id}. */
  name?: string
  /** Advisory models exposed to replay scenarios that exercise discovery. */
  models?: ReplayModelConfig[]
  /** Optional provider-owned retry policy used by assembled recovery snapshots. */
  retryPolicy?: RetryPolicyConfig
}

/** One model exposed by a replay-only provider catalog. */
export interface ReplayModelConfig {
  /** Model id used for replay requests. */
  id: string
  /** Selector label; defaults to {@link id}. */
  name?: string
  /** Optional selector description. */
  description?: string
  /** Optional positive integer context capacity published by the replay adapter. */
  contextWindow?: number
  /** Optional declared input modalities, so a scenario can exercise capability gates (e.g. image-capable
  `read_image`). */
  inputModalities?: readonly ModelModality[]
}

```



```

* Optional per-request output cap the replay route materializes when callers
* omit one, so replay reconstructs the request header a live catalog produced.
*/
defaultMaxTokens?: number
/**
* Optional flat visual-token price the replay route declares for every
* retained request image, so keyless scenarios exercise route-priced
* request pressure; each occurrence is priced at this value plus its
* request-preview handle text. Requires {@link inputModalities} to include
* `image` — a text-only route never sends visual tokens. Absent declares
* no image pricing.
*/
imageRequestTokens?: number
/** Optional reasoning-effort ids the replay route accepts, in display order. */
reasoningEfforts?: string[]
/**
* Optional effort materialized when callers omit one; must appear in
* {@link reasoningEfforts} or call resolution rejects the route.
*/
defaultReasoningEffort?: string
}

```

[ModelModality](#) · [RetryPolicyConfig](#)

[packages/test-support/llm-replay/src/index.ts:924](#)

@deepseek-ai/dsh-llm-retry

[agents](#) · [sessionProjections](#)

```

/** This policy executor has no config; providers own `retryPolicy`. */
export type Config = Readonly<Record<string, never>>

```

[packages/llm/llm-retry/src/index.ts:25](#)

@deepseek-ai/dsh-lsp-stdio

```

/** Plugin configuration: provider id → local language-server configuration. */
export interface Config {
  /** Non-empty table of stable provider ids to independent local server configurations. */
  servers: Record<string, LspLocalServerConfig>
}

/** One configured local language server and its host bounds. */
export interface LspLocalServerConfig {
  /** Executable to spawn (absolute, or resolved on PATH at load). */
  command: string
  /** Lowercase leading-dot extension → LSP language id (e.g. `{ '.ts': 'typescript' }`). */
  extensionToLanguage: Record<string, string>
  /** Arguments passed to the executable (no shell). Default `[]`. */
  args?: string[]
  /** Extra env vars merged on top of the scrubbed ambient env. Default `{}`. */
  env?: Record<string, string>
  /** Static `initialize` options forwarded to the server. Default `null`. */
  initializationOptions?: unknown
  /** Static answer to every `workspace/configuration` item. Default `null`. */
  configuration?: unknown
  /** Largest single framed message accepted from the server (bytes). Default 16000000. */
  maxMessageBytes?: number
  /** Largest stderr tail retained for diagnostics (bytes). Default 1000000. */
  maxStderrBytes?: number
  /** Largest source file this host will open (bytes). Default 4000000. */
  maxDocumentBytes?: number
  /** Graceful `shutdown`/`exit` budget before escalation (ms). Default 5000. */
  shutdownTimeoutMs?: number
  /** Request-cancel and SIGTERM→SIGKILL grace (ms). Default 2000. */
  killGraceMs?: number
}

```

[packages/lsp/lsp-stdio/src/index.ts:82](#)

@deepseek-ai/dsh-mcp-client

```

/** Configuration for one stdio or Streamable HTTP MCP server. */
export type Config = StdioConfig | StreamableHttpConfig

/** Config for connecting to an MCP server via a spawned child process over stdio. */
export interface StdioConfig {
  /** Selects child-process stdio transport. */
  transport: 'stdio'
  /**
   * Stable local namespace for this server's model-facing tool names
   * (`mcp__<serverName>__<rawName>`). Must match `[A-Za-z0-9_-]{1,32}` and be
   * unique across live mcp-client instances.
   */
  serverName: string
  /** Executable used to start the server. */
  command: string
  /** Arguments passed directly, without shell interpolation. */
  args: string[]
  /** Extra env vars merged on top of scrubbed ambient env. */
  env: Record<string, string>
  /** Working directory for the child process. */
  cwd: string
  /** Per-tool-call timeout in milliseconds. */
  toolCallTimeoutMs: number
  /** Fail plugin activation when the initial connection or tool synchronization fails. */
  failOnStartupError: boolean
  /** Automatic reconnect policy after a lost connection; omission uses the defaults. */
  reconnect?: ReconnectConfig
}

/** Config for connecting to an MCP server over Streamable HTTP (SSE). */
export interface StreamableHttpConfig {
  /** Selects Streamable HTTP transport. */
  transport: 'streamable-http'
  /**
   * Stable local namespace for this server's model-facing tool names
   * (`mcp__<serverName>__<rawName>`). Must match `[A-Za-z0-9_-]{1,32}` and be
   * unique across live mcp-client instances.
   */
  serverName: string
  /**
   * Stable local namespace for this server's model-facing tool names
   * (`mcp__<serverName>__<rawName>`). Must match `[A-Za-z0-9_-]{1,32}` and be
   * unique across live mcp-client instances.
   */
}

```

```

serverName: string
/** MCP endpoint URL. */
url: string
/** Additional headers attached to MCP requests. */
headers: Record<string, string>
/** Per-tool-call timeout in milliseconds. */
toolCallTimeoutMs: number
/** Fail plugin activation when the initial connection or tool synchronization fails. */
failOnStartupError: boolean
/** Automatic reconnect policy after a lost connection; omission uses the defaults. */
reconnect?: ReconnectConfig
}

/** Automatic reconnect policy for one MCP server connection. */
export interface ReconnectConfig {
  /** Reconnect automatically after a lost connection (default true). */
  enabled?: boolean
  /** First reconnect delay in milliseconds; doubles per consecutive failed attempt (default 500). */
  initialDelayMs?: number
  /** Backoff ceiling in milliseconds; also the uptime after which the attempt budget resets (default 30000). */
  maxDelayMs?: number
  /** Consecutive failed attempts per outage before giving up for good (default 10). */
  maxAttempts?: number
}

```

[packages/mcp/mcp-client/src/index.ts:98](#)

@deepseek-ai/dsh-message-feedback

[storageDomain](#) · [sessionPersistence](#) · [sessions](#)

```

/** Required deployment policy for optional notes. */
export interface Config {
  /** Maximum UTF-8 byte length accepted for one note. */
  readonly maxNoteBytes: number
}

```

[packages/feedback/message-feedback/src/index.ts:50](#)

@deepseek-ai/dsh-permission-presets

shell · approval · sessions · sessionProjections

```
/** The {@link PermissionPresetService} config: preset table and composition default. */
export interface Config {
  /**
   * The preset table: name → knob bundle. Defaults to `workspace-write`
   * (workspace-write + ask) and `danger-full-access` (danger-full-access +
   * never). The name `custom` is reserved for the derived not-a-preset state.
   */
  presets?: Record<string, PresetSpec>
  /**
   * Default for new sessions. When omitted, the preset matching the composed
   * sandbox and approval defaults is used.
   */
  defaultPreset?: string
}

/** One preset's sandbox/approval bundle and optional client presentation. */
export interface PresetSpec {
  /** The `sandbox/mode` value the preset writes through. */
  sandbox: SandboxMode
  /** The `approval/policy` value the preset writes through. */
  approval: ApprovalPolicy
  /** The display label a client shows for this preset; the raw table key when omitted. */
  name?: string
  /** One user-facing sentence on what the preset means; omitted when not configured. */
  description?: string
}
```

[ApprovalPolicy](#) · [SandboxMode](#)

[packages/interaction/permission-presets/src/index.ts:143](#)

@deepseek-ai/dsh-persona

systemPrompt

```
/** Plugin config: the persona text this composition contributes. */
export interface Config {
  /**
   * Persona prose rendered as the `deployment:persona` section. A template:
   * complete `${...}` groups interpolate strictly against registered prompt
   * variables. Empty text drops the section at render, matching the registry.
   */
  text: string
  /** Make this persona the complete system prompt, suppressing every other section. */
  complete?: boolean
  /** Suppress dynamic runtime-context snapshots for this persona's agent scope. */
  includeRuntimeContext?: boolean
}
```

[packages/preset/persona/src/index.ts:30](#)

@deepseek-ai/dsh-plan-mode

tools · systemPrompt · sessionProjections

```
/** Deployment-owned plan guidance. */
export interface PlanModeConfig {
  /** Guidance rendered as the `plan:policy` prompt section while plan mode is active. */
  section: string
}
```

[packages/plan/plan-mode/src/index.ts:63](#)

@deepseek-ai/dsh-plugin-package-inventory-deepseek

agents · deepseekLlmApiExtensions · loader

```
/** Plugin-package request contribution configuration. */  
export interface Config {  
  /** Contribute `dsh_plugin_packages` to official DeepSeek requests. Defaults to `true`. */  
  enabled?: boolean  
}
```

[packages/llm/plugin-package-inventory-deepseek/src/index.ts:31](#)

@deepseek-ai/dsh-pwsh-local

subprocess

```
/** Plugin config (all optional — `static Config` supplies the defaults). */  
export interface Config {  
  /** Default working directory for commands (default: process.cwd()). */  
  cwd?: string  
  /** Default foreground timeout in milliseconds. */  
  timeoutMs?: number  
  /** Upper bound for per-call timeout overrides. */  
  maxTimeoutMs?: number  
  /** Per-stream in-memory output cap; overflow spills to a temp file. */  
  maxOutputBytes?: number  
  /** Per-stream spill-file cap; larger streams retain only their in-memory tail. */  
  maxSpillBytes?: number  
  /** Grace period for kill escalation and inherited pipes; at most `MAX_TIMER_DELAY_MS`. */  
  graceMs?: number  
  /**  
   * Explicit pwsh executable. When omitted, well-known Windows install  
   * locations and PATH entries are probed in order (PowerShell 7 install,  
   * PATH entries such as the Microsoft Store install, then Windows  
   * PowerShell 5.1), falling back to a bare `pwsh` resolved through PATH.  
   */  
  pwshPath?: string  
}
```

[packages/shell/pwsh-local/src/index.ts:58](#)

@deepseek-ai/dsh-pwsh-sandbox

subprocess · sandbox · sandboxPolicy

```
/**
 * Plugin config: the local executor's knobs, verbatim. The sandbox policy —
 * the default mode and fallback `workspace-write` root — is NOT here: it lives
 * on `ctx.sandboxPolicy` (`@deepseek-ai/dsh-sandbox-policy`), which resolves
 * each calling session's mode and cwd for every enforcing capability. The
 * runner choice is likewise the `ctx.sandbox` provider's config, not this
 * executor's.
 */
export type Config = LocalConfig
```

[LocalConfig](#)

[packages/shell/pwsh-sandbox/src/index.ts:40](#)

@deepseek-ai/dsh-repeat-tool-reminder

```
/**
 * Plugin config, validated by the same-named schemastery schema plus the
 * load-time checks in `apply` (misconfiguration fails loud: an empty
 * `thresholds` list, a non-integer, a value below 2, or a duplicate throws at
 * plugin load, never a silent fall-back). `include`/`exclude` entries are
 * `*-wildcard predicates over tool names at call time, not references to
 * registry entries — a pattern matching no currently registered tool is valid
 * (`exclude: [mcp_*]` must stay legal in a deployment that loads no MCP tools).
 */
export interface Config {
  /** Consecutive-repeat counts that trigger a reminder (default `[3, 5, 8]`). */
  thresholds?: number[]

  /** Tool-name patterns to track; empty means every tool is tracked. */
  include?: string[]

  /** Tool-name patterns transparent to the chain (neither count nor reset). */
  exclude?: string[]
}

/**
 * Maximum characters of canonical arguments quoted in the DETAILED reminder
```



```

* (default 500). Large payloads (a `write` body, a long command) would
* otherwise ride into the next request unbounded — precisely in a loop
* scenario; the cap bounds the reminder, never the detection (the chain key
* always compares the FULL canonical string).
*/
argumentsPreviewChars?: number
}

```

[packages/guard/repeat-tool-reminder/src/index.ts:28](#)

@deepseek-ai/dsh-sandbox-local

```

/** Plugin config. All optional — `static Config` supplies the defaults. */
export interface Config {
  /**
   * Override the runner argv; bwrap-compatible profile arguments are appended. A
   * non-empty override asserts full enforcement and skips built-in selection and
   * probing. A runner that starts but refuses its profile must be identifiable by
   * {@link runnerFailureSignatures}. Consumers classify a spawn rejection only after
   * confirming the workdir is usable. `ENOENT` or `EACCES` identifies the runner when
   * `error.path` equals argv[0] and `error.syscall` is `spawn` or `spawn <runner>`, or
   * when `error.path` is absent and `error.syscall` is exactly `spawn <runner>`.
   */
  runnerCommand?: string[]
  /**
   * Case-insensitive stderr substrings emitted when a configured
   * {@link runnerCommand} refuses its profile before executing the wrapped
   * command. Required and non-empty with `runnerCommand`; rejected without
   * it. Each entry is a non-empty, single-line, case-insensitive substring
   * covering the executable runner's own failure dialect.
   */
  runnerFailureSignatures?: string[]
  /** Positive timeout for each functional probe; zero would mean unbounded to Node. */
  probeTimeoutMs?: number
}

```

[packages/sandbox/sandbox-local/src/index.ts:44](#)

@deepseek-ai/dsh-sandbox-policy

sessionProjections

```
/**
 * Plugin config: the deployment's sandbox default. All optional — `Config`
 * supplies the defaults (`mode: 'read-only'` is the fail-safe default; a
 * deployment that wants a workspace-writable agent opts in explicitly). The
 * runner choice is NOT here (it is the `ctx.sandbox` provider's config), nor
 * is any per-family knob: this is the one shared policy home.
 */
export interface Config {
  /** File-sandbox mode a session starts from (default: `read-only`). */
  mode?: SandboxMode
  /**
   * Fallback root for agentless calls and sessions without a cwd (default:
   * `process.cwd()`). Normal agent calls use their session cwd instead.
   */
  workspaceRoot?: string
}
```

[SandboxMode](#)

[packages/sandbox/sandbox-policy/src/index.ts:70](#)

@deepseek-ai/dsh-sdk-app

cmdlineArgs

```
/** SDK stdio startup configuration. */
export interface Config {
  /** Profile name rendered in help and diagnostics (default `sdk`). */
  profile?: string
}
```

[packages/bundle/sdk-app/src/index.ts:23](#)

@deepseek-ai/dsh-sdk-jsonrpc-server

agents

```
/** JSON-RPC deployment config plus runtime-only test hooks. */
export interface JsonRpcConfig {
  /** Report max-token turn/subagent termination as a successful SDK result. */
  maxTokensAsSuccess?: boolean
  /** Transport input override; production uses `process.stdin`. */
  input?: Readable
  /** Transport output override; production uses `process.stdout`. */
  output?: Writable
  /** Process-exit override; production uses `process.exit`. */
  exit?: (code: number) => void
}
```

Readable node:stream ☐ Writable node:stream ☐

[packages/sdk/server/src/index.ts:25](#)

@deepseek-ai/dsh-session-log-deepseek

deepseekLlmApiExtensions · sessions

```
/** Session-log request contribution configuration. */
export interface Config {
  /** Contribute `dsh_session_log` to official DeepSeek requests. Defaults to `false`. */
  enabled?: boolean
}
```

[packages/session/session-log-deepseek/src/index.ts:36](#)

@deepseek-ai/dsh-session-log-export

commands · connection

```

/** Session-log archive policy. */
export interface Config {
  /** DEFLATE level for each ZIP entry. @default 6 */
  readonly compressionLevel?: SessionLogCompressionLevel
}

/** Valid fflate DEFLATE levels accepted by session-log export. */
export type SessionLogCompressionLevel = 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9

```

[packages/session-query/session-log-export/src/index.ts:42](https://github.com/deepseek-ai/dsh-session-persistence-jsonl/blob/main/packages/session-query/session-log-export/src/index.ts#L42)

@deepseek-ai/dsh-session-persistence-jsonl

sessions · sessionProjections

```

/** Plugin config: where the JSONL backend keeps its session logs, and the packed-row write switch. */
export interface Config {
  /**
   * Root directory for all session files. Required (no default): a default of
   * `process.cwd()` would scatter session files as the process's cwd changes
   * (bash calls, subprocesses). Sessions group under human-readable project
   * directories, then per-session directories. An existing root must be a
   * readable directory; an absent root is created on first materialization.
   */
  root: string
  /**
   * Write runs of consecutive `assistant/chunk` delta events as packed
   * `text-chunks`/`reasoning-chunks`/`tool-call-chunks` rows (lossless,
   * ~60% smaller logs measured on a real session). Defaults to true; false
   * keeps one `SessionEvent` per line for diagnostics. Reading packed rows is
   * unconditional: a log's layout never depends on this switch.
   */
  packChunks?: boolean
  /** Physical encoding; defaults to checksummed Zstandard frames. */
  compression?: JsonlCompression
  /** Maximum cold Session preparations retained for history-to-resume reuse. */

```

```
preparedSessionCacheSize?: number

/** Fixed live-event coalescing window; not a backend completion deadline. */

writeBatchMaxDelayMs?: number
}

/** Physical encoding selected for JSONL session artifacts. */
export type JsonlCompression = 'zstd' | 'none'
```

[packages/session/session-persistence-jsonl/src/index.ts:70](#)

@deepseek-ai/dsh-session-projection-cache

storageDomain · sessionProjections · sessions

```
/**
 * Plugin config. Both throttle triggers are deployment choices with no
 * universally correct value, so the composition states them explicitly
 * (cordis.yml); the three mandatory write points (session creation,
 * `turn/end`, and session disposal) are policy, not tunables, and always
 * fire.
 */
export interface Config {
  /** Committed events per session that force a durable checkpoint write between mandatory points. */
  writeEveryEvents: number

  /** Longest time (milliseconds) a dirty checkpoint may stay unwritten between mandatory points. */
  writeIntervalMs: number
}
```

[packages/session/session-projection-cache/src/index.ts:55](#)

@deepseek-ai/dsh-session-query-sqlite

sessions

```

/** Combined session-query configuration backed by SQLite full-text search. */
export interface Config extends SessionQueryConfig {
  /**
   * Dedicated derived-index path; `:memory:` is supported for ephemeral
   * indexes. Missing directories and database files are created owner-only on
   * POSIX filesystems; existing modes are preserved.
   */
  path: string
  /**
   * Open the SQLite module and handle at service activation or the first
   * search, or `never` to disable full-text search: the inherited exact
   * reads, filters, and traces stay available, while `searchSessions` and
   * `searchEvents` fail with `SESSION_QUERY_SEARCH_DISABLED` and SQLite is
   * never imported or opened. Defaults to `startup`.
   */
  openAt?: OpenAt
  /** SQLite journal mode. Defaults to `wal`. */
  journalMode?: JournalMode
  /** Page size when a request omits `limit`. At most `Number.MAX_SAFE_INTEGER - 1`; defaults to 20. */
  defaultLimit?: number
  /** Largest accepted page size. At most `Number.MAX_SAFE_INTEGER - 1`; defaults to 100. */
  maxLimit?: number
  /** Maximum snippet length in Unicode code points. Defaults to 240. */
  snippetChars?: number
  /** Maximum concurrent persisted-log inspections in one inherited batch read. Defaults to 4. */
  persistedInspectConcurrency?: number
}

/** SQLite module/handle opening phase; `never` disables full-text search entirely. */
export type OpenAt = 'startup' | 'first-search' | 'never'

/** Supported SQLite journal modes. */
export type JournalMode = 'wal' | 'delete' | 'truncate' | 'persist'

```

[SessionQueryConfig](#)

[packages/session-query/session-query-sqlite/src/index.ts:96](#)

@deepseek-ai/dsh-session-reference

sessionQuery

```
/** Session-reference service configuration. */  
export interface Config {  
  /** Maximum distinct source sessions referenced by one message, from one to three. */  
  maxReferences?: number  
  /** Default host candidate-list limit. */  
  candidateLimit?: number  
  /** Maximum rendered UTF-8 bytes for one source snapshot. */  
  maxReferenceBytes?: number  
}
```

[packages/context/session-reference/src/config.ts:11](#)

@deepseek-ai/dsh-session-telemetry-otel

sessions

```
/**  
 * Plugin configuration: one sharing policy, two verbatim SDK option objects,  
 * and one DSH-owned shutdown bound. Uploading modes validate their endpoint  
 * and shutdown deadline at plugin load; `DISABLED` reads neither.  
 */  
export interface Config {  
  /** Sharing policy; defaults to local-only `DISABLED` behavior. */  
  mode?: SessionTelemetryMode  
  /**  
   * Passed verbatim to the SDK's OTLP/HTTP log exporter — the complete  
   * `OTLPExporterNodeConfigBase` shape (`headers`, `timeoutMillis`,  
   * `compression`, `keepAlive`, ...), owned and documented by the SDK. `url`  
   * is the one field this package requires and validates itself.  
   */  
  exporter?: OTLPExporterNodeConfigBase & {  
    /** Full logs endpoint (e.g. `https://collector.example.com/v1/logs`). Required outside `DISABLED`; validated  
    at load. */  
    url?: string
```

```

}
/**
 * Passed verbatim to `BatchLogRecordProcessor` (minus the exporter slot,
 * which this plugin fills); the SDK owns and documents these knobs.
 */
processor?: Omit<BatchLogRecordProcessorOptions, 'exporter'>
/** Maximum time spent awaiting the SDK provider's complete shutdown path. */
shutdownTimeoutMillis?: number
}

/** Session-sharing policy selected by {@link Config.mode}. */
export enum SessionTelemetryMode {
  FULL = 'FULL',
  FEEDBACK_ONLY = 'FEEDBACK_ONLY',
  DISABLED = 'DISABLED',
}

```

BatchLogRecordProcessorOptions @opentelemetry/sdk-logs  OTLPExporterNodeConfigBase @opentelemetry/otlp-exporter-base 

[packages/session/session-telemetry-otel/src/index.ts:91](#)

@deepseek-ai/dsh-session-title

sessions

```

/** Required deterministic fallback and accepted-title limits. */
export interface Config {
  /** Maximum whitespace-delimited words in the built-in fallback. */
  readonly fallbackMaxWords: number
  /** Maximum UTF-8 bytes in the built-in fallback. */
  readonly fallbackMaxBytes: number
  /** Maximum UTF-8 bytes in any accepted title. */
  readonly maxTitleBytes: number
}

```

[packages/session/session-title/src/index.ts:56](#)

@deepseek-ai/dsh-session-title-all-prompts-llm

sessionTitle · llm · sessions

```
/** Required LLM policy; this plugin adds no defaults. */  
export type Config = SessionTitleLlmConfig
```

[SessionTitleLlmConfig](#)

[packages/session/session-title-all-prompts-llm/src/index.ts:15](#)

@deepseek-ai/dsh-session-title-first-prompt-llm

sessionTitle · llm · sessions

```
/** Required LLM policy; this plugin adds no defaults. */  
export type Config = SessionTitleLlmConfig
```

[SessionTitleLlmConfig](#)

[packages/session/session-title-first-prompt-llm/src/index.ts:15](#)

@deepseek-ai/dsh-settings-file

```
/** Plugin config: file location and hot-reload behavior. */  
export interface Config {  
  /** Settings document path; defaults to `settings.yaml` under the harness home. */  
  path?: string  
  /** Harness home used when `path` is omitted; defaults to `${DSH_HOME}` or `~/.dsh`. */  
  dshHome?: string  
  /** Watch the document and hot-publish external edits; defaults to true. */  
  watch?: boolean
```

```
/** Watcher write-settle window in milliseconds; defaults to 100. */
debounceMs?: number
}
```

[packages/settings/settings-file/src/index.ts:22](#)

@deepseek-ai/dsh-shell-env

```
/** Plugin config (all optional — the built-in facts resolve without defaults). */
export interface Config {
  /** DeepSeek Harness home directory exposed as `DSH_HOME`; defaults to `$DSH_HOME` or `~/.dsh`. */
  dshHome?: string
}
```

[packages/shell/shell-env/src/index.ts:29](#)

@deepseek-ai/dsh-skill

```
/** Skill registry configuration. */
export interface Config {
  /** Maximum number of completed cwd/provider catalogs kept in memory. */
  readonly collectCacheMaxEntries?: number
}
```

[packages/skill/skill/src/index.ts:280](#)

@deepseek-ai/dsh-skill-filesystem

skills

```
/** Local filesystem skill provider configuration. */
export interface Config {
  /** Unique provider name. Defaults to `filesystem`. */
  providerName?: string
}
```

```

/** Whether project and user roots are included around custom roots. */
includeDefaultRoots?: boolean

/** DeepSeek Harness config root. Defaults to ` $DSH_HOME ` or ` ~/.dsh `. */
dshHome?: string

/** Shared agent config root. Defaults to ` $DSH_AGENTS_HOME ` or ` ~/.agents `. */
agentsHome?: string

/** Additional skill roots scanned after project roots and before user roots. */
customSkillDirs?: string[]

/** Whether host-local skill roots are watched for catalog changes. */
watch?: boolean

/** Whether Chokidar uses polling instead of native filesystem events. */
watchUsePolling?: boolean

/** Milliseconds a changed skill entry must remain stable before it is observed. */
watchStabilityThresholdMs?: number

/** Milliseconds between Chokidar stability or polling probes. */
watchPollIntervalMs?: number

/** Maximum distinct project roots whose skill directories remain watched. */
watchMaxProjects?: number

/** Whether watched symbolic links follow their target files. */
watchFollowSymlinks?: boolean

/** Bundled skill root; defaults to ` $DSH_BUNDLED_SKILL_DIR ` when default roots are included, otherwise
mounts none. */
bundledSkillDir?: string
}

```

[packages/skill/skill-file-system/src/index.ts:49](#)

@deepseek-ai/dsh-spill-local

```

/** Plugin config (all optional — `static Config` supplies the defaults). */
export interface Config {
  /**
   * Root directory for spill files. Omitted uses a lazily-created private
   * (0700) per-process directory under the OS temp dir — the safe default for
   * a local deployment. Set it to keep spill files under a known location.
   */
  root?: string
}

```

```
* Age in days after which a spill file is eligible for the one-shot startup
* cleanup sweep. Defaults to `30`; `0` disables cleanup entirely. Files whose
* `mtime` is strictly older than the cutoff are deleted and emptied
* directories are pruned; fresh files, symlinks, and unrelated entries are
* left untouched. On POSIX, cleanup skips roots and session directories that
* another local user could modify or replace. Retention is deliberate — a
* resumed or forked session may still reference an older locator until it
* ages out.
*/
cleanupPeriodDays?: number
}
```

[packages/spill/spill-local/src/index.ts:31](#)

@deepseek-ai/dsh-spill-policy

tools · sessionProjections

```
/** Plugin config. */
export interface Config {
  /**
   * The model-facing context cap for a plain-text tool result, in UTF-8 bytes.
   * Omitted disables the policy entirely (no-op). When set, a result larger than
   * this is spilled and replaced with a preview derived from this same budget.
   */
  maxInlineBytes?: number
}
```

[packages/spill/spill-policy/src/index.ts:60](#)

@deepseek-ai/dsh-storage-domain

storage

```
/**
 * Plugin config. Which backend serves which domain is decided here, not
 * globally on the hub: `backend` is the default route and `routes` overrides
```

```
* it per domain name. A route naming an unregistered backend fails loud at
* `open` with `backend-not-found`.
*/
export interface Config {
  /** Default backend name for every domain without an explicit route. Required: there is no universally correct
  medium. */
  backend: string
  /** Per-domain overrides: domain name → backend name. */
  routes?: Record<string, string>
}
```

[packages/storage/storage-domain/src/index.ts:52](#)

@deepseek-ai/dsh-storage-json

storage

```
/**
 * Plugin configuration.
 * `root` has NO default on purpose: a `process.cwd()` fallback would scatter
 * unit files wherever the process happens to start; assemblies state the
 * location explicitly.
 */
export interface Config {
  /** Directory holding one `.json` file (or `` tree) per unit. */
  root: string
}
```

[packages/storage/storage-json/src/index.ts:28](#)

@deepseek-ai/dsh-storage-sqlite

storage

```
/** Plugin configuration. */
export interface Config {
  /**
```

- * Filesystem path to the SQLite database file. The special value `:memory:`
- * opens an in-process database (tests). On filesystems with POSIX modes,
- * missing directories and databases are created owner-only; existing path
- * modes are preserved. Filesystem setup errors other than an existing
- * database fail the open. The backend does not protect confidentiality or
- * integrity when another principal can replace the database entry in its
- * parent directory.

*/

path: string

/**

- * SQLite `journal\_mode` pragma. `wal` (the default) suits local disks; pick
- * a rollback-journal mode (`delete`/`truncate`/`persist`) on filesystems
- * where WAL's shared-memory files do not work (network mounts). See
- * {@link JournalMode}.

*/

journalMode?: JournalMode

}

/**

- * Journal modes the backend will run under. `wal` is the default; the
- * rollback-journal modes (`delete`/`truncate`/`persist`) exist for
- * filesystems where WAL's shared-memory files do not work (network mounts).
- * `memory`/`off` are excluded: dropping journal durability silently
- * contradicts the durability clause of the KV backend contract.

*/

export type JournalMode = 'wal' | 'delete' | 'truncate' | 'persist'

[packages/storage/storage-sqlite/src/index.ts:24](#)

@deepseek-ai/dsh-subagent-acp

subagents · subprocess

/** Config: how to spawn and drive the child ACP agent process. */

export interface Config {

/** Provider name on `ctx.subagents` (default `acp`). */

providerName: string

/** The executable to spawn for each run (the child ACP agent). */

```

command: string
/** Arguments passed to {@link command}. */
args: string[]
/**
 * Working directory override for the child process and its ACP session.
 * Must be non-empty; a relative path resolves against the harness launch
 * directory at load, and the result must be an existing directory. When
 * omitted, each child inherits its delegating parent session's cwd — and
 * starting one from a parent session that has no cwd fails.
 */
cwd?: string
/**
 * How to auto-answer the child's `session/request_permission` prompts:
 * `reject` (default — decline every prompt) or `allow` (approve via the first
 * `allow_once` or `allow_always` option). No prompt is surfaced to a human.
 */
permission: PermissionPolicy
/**
 * Extra environment variables for the child process — e.g. the child
 * harness's own `DEEPSEEK_API_KEY`. Forwarded on top of a credential-scrubbed
 * copy of the parent env, so an explicit key here reaches the child while
 * ambient secrets do not leak implicitly.
 */
env: Record<string, string>
/**
 * Grace period (ms) for the child's EOF-driven quiesce on dispose — its
 * window to flush persistence and tear down its own nested subprocesses
 * before the parent escalates to a signal. Must not exceed
 * `MAX_TIMER_DELAY_MS`.
 */
disposeEofGraceMs?: number
/** Failure-observation and termination-escalation grace (ms); must not exceed `MAX_TIMER_DELAY_MS`. */
disposeGraceMs?: number
}

/** Fixed response to child permission requests: reject by default, or select the first allow option. */
export type PermissionPolicy = 'allow' | 'reject'

```

@deepseek-ai/dsh-subagent-claude-code

subagents · subprocess

```
/** Deployment-owned model, permission, environment, and process-release settings. */
export interface Config {
  /** Provider name on `ctx.subagents` (default `claude-code`). */
  providerName?: string
  /** Native Claude model fixed for this instance; omitted to inherit Claude settings. */
  model?: string
  /**
   * Explicit environment entries layered over the subprocess seam's
   * credential-scrubbed parent environment.
   */
  env?: Record<string, string>
  /**
   * Native non-interactive mode fixed for this Provider instance. Defaults to
   * `dontAsk`; `acceptEdits` accepts edits, `auto` uses the native classifier,
   * `plan` returns a plan without approving execution, and
   * `bypassPermissions` explicitly skips permission checks.
   */
  permissionMode?: ClaudeCodePermissionMode
  /** Grace in milliseconds for Claude Code process-tree termination. */
  disposeGraceMs?: number
}

/** Profile-selectable non-interactive Claude Code permission mode. */
export type ClaudeCodePermissionMode = typeof CLAUDE_CODE_PERMISSION_MODES[number]
```

[packages/subagent/subagent-claude-code/src/index.ts:38](https://github.com/deepseek-ai/dsh-subagent-claude-code/blob/main/packages/subagent/subagent-claude-code/src/index.ts#L38)

@deepseek-ai/dsh-subagent-codex

subagents · subprocess


```

/** Deployment-owned model, permission, environment, and process-release settings. */
export interface Config {
  /** Provider name on `ctx.subagents` (default `codex`). */
  providerName?: string
  /** Native Codex model fixed for this instance; omitted to inherit Codex settings. */
  model?: string
  /**
   * Explicit environment entries layered over the subprocess seam's
   * credential-scrubbed parent environment.
   */
  env?: Record<string, string>
  /** Native non-interactive permission mode fixed for this Provider instance. */
  permissionMode?: CodexPermissionMode
  /** Grace in milliseconds for app-server process-tree termination. */
  disposeGraceMs?: number
}

```

```

/** Profile-selectable non-interactive Codex permission mode. */
export type CodexPermissionMode =
  | 'never'
  | 'approve-for-me'
  | 'dangerously-bypass-approvals-and-sandbox'

```

[packages/subagent/subagent-codex/src/index.ts:36](https://github.com/deepseek-ai/dsh-subagent-dsh-sdk/blob/main/packages/subagent/subagent-codex/src/index.ts#L36)

@deepseek-ai/dsh-subagent-dsh-sdk

subagents

```

/** Config: how to spawn and drive the child SDK runtime process. */
export interface Config {
  /** Provider name on `ctx.subagents` (default `dsh-sdk`). */
  providerName: string
  /** Explicit dsh CLI module, resolved and checked at plugin load; omission uses the SDK dependency. */
  dshBin?: string
  /** Named child profile (default `sdk`). */
  profile: string
  /** Ordered per-launch profile patch files, resolved and checked at plugin load. */

```

```

patches: string[]
/** Absolute isolated Harness home for every nested child process. */
dshHome: string
/**
 * Working directory override for the child process and its SDK session
 * workspace. Must be non-empty; a relative path resolves against the
 * harness launch directory at load, and the result must be an existing
 * directory. When omitted, each child inherits its delegating parent
 * session's cwd — and starting one from a parent session that has no cwd
 * fails.
 */
cwd?: string
/** Provider route the child runtime initializes with (default `deepseek-official`). */
provider: string
/** Model the child runtime initializes with (default `deepseek-v4-flash`). */
model: string
/** Optional per-request output-token cap for the child runtime. */
maxTokens?: number
/**
 * Extra environment variables for the child process — e.g. the child
 * runtime's own `DEEPSEEK_API_KEY`. Forwarded on top of a credential-scrubbed copy of the parent
 * env, so an explicit key here reaches the child while ambient secrets do
 * not leak implicitly.
 */
env: Record<string, string>
/** Bound (ms) on the protocol `shutdown` exchange during dispose. */
shutdownTimeoutMs?: number
/**
 * Grace period (ms) for the child's EOF-driven quiesce on dispose — its
 * window to flush persistence and tear down its own nested subprocesses
 * before the parent escalates to a signal.
 */
disposeEofGraceMs?: number
/** Termination confirmation window (ms), including forced exit on every platform. */
disposeGraceMs?: number
}

```

@deepseek-ai/dsh-subagent-fork-in-process

subagents

```
/** Config: the registry name to register the provider under. */  
export interface Config {  
  /** Provider name on `ctx.subagents` (default `fork`). */  
  providerName: string  
}
```

[packages/subagent/subagent-fork-in-process/src/index.ts:31](#)

@deepseek-ai/dsh-subagent-spawn-in-process

subagents

```
/** Config: the registry name to register the provider under. */  
export interface Config {  
  /** Provider name on `ctx.subagents` (default `spawn`). */  
  providerName: string  
}
```

[packages/subagent/subagent-spawn-in-process/src/index.ts:25](#)

@deepseek-ai/dsh-subprocess-e2b

e2b

```
/** Configuration for the E2B subprocess adapter. */  
export interface Config {  
  /** Remote status/liveness poll cadence in milliseconds; each tick is one control-plane request. */  
  pollMs?: number  
}
```

```
}
```

[packages/e2b/subprocess-e2b/src/index.ts:25](#)

@deepseek-ai/dsh-system-prompt

```
/** Plugin config: the deployment-authored fragment of the system prompt (see {@link Config.persona} for its
contract). */
export interface Config {
  /** Include the fixed DeepSeek Harness identity before the deployment persona (default true). */
  includeHarnessIdentity?: boolean
  /** Include dynamic runtime-context snapshots in model history (default true). */
  includeRuntimeContext?: boolean
  /**
   * Deployment-wide order-0 persona template. A scoped section named
   * `deployment:persona` shadows it; `{{variable}}` references are strict.
   */
  persona?: string
  /**
   * Model-facing tool names in order, with {@link TOOL_ORDER_REST} exactly once.
   * Invalid fields fail at load and unknown names fail at assembly; known names
   * hidden in one scope may be absent there. Omitted means lexicographic order.
   */
  toolOrder?: string[]
}
```

[packages/core/system-prompt/src/index.ts:237](#)

@deepseek-ai/dsh-terminal-bash

[terminals](#) · [sandboxPolicy](#) · [sessionProjections](#) · [subprocess](#)

```
/** Public plugin configuration. */
export interface Config {
  /** Backend registry type (default: `shell`). */
  backendType?: string
}
```

```

/** Interactive shell dialect (default: `bash`); selects the argv/env/startup defaults. */
shellDialect?: ShellDialect

/** Interactive shell executable (default per dialect: `/bin/bash`, or the resolved pwsh). */
shellPath?: string

/** Shell arguments (default per dialect: bash `--noprofile --norc -i`, pwsh `-NoLogo -NoProfile`). */
shellArgs?: string[]

/** Terminal rows. */
rows?: number

/** Terminal columns. */
cols?: number

/** Maximum retained logical lines. */
scrollbackLines?: number

/** Maximum retained UTF-8 bytes. */
scrollbackMaxBytes?: number

/** Maximum bytes returned by one read or settled viewport. */
maxReadBytes?: number

/** Readiness polling interval. */
pollIntervalMs?: number

/** Delay before Linux exact syscall probes. */
exactProbeAfterMs?: number

/** Silence duration that yields `inferred_idle`. */
idleSilenceMs?: number

/**
 * Extra wait beyond `idleSilenceMs`, once a prompt marker was seen, for the shell to
 * regain the foreground before `inferred_idle` settles; at least one `pollIntervalMs`.
 */
handoffGraceMs?: number

/** Absolute bound for one send and the complete pwsh startup sequence. */
timeoutMs?: number

/** Grace before teardown escalates to `SIGKILL`. */
disposeGraceMs?: number
}

/** One supported interactive shell dialect. */
export type ShellDialect = 'bash' | 'pwsh'

```

[packages/terminal/terminal-bash/src/config.ts:10](https://github.com/microsoft/terminal/blob/main/packages/terminal/terminal-bash/src/config.ts#L10)

@deepseek-ai/dsh-time-context

agents · sessionProjections

```
/** Request-preparation clock formatting and append scheduling. Invalid values fail plugin load. */
export interface Config {
  /** Fallback display zone when the open turn has no unique browser zone. Omit to use the process zone. */
  timeZone?: string
  /** Minimum milliseconds between durable injections in one session. Omit or set to 0 to inject at every eligible step. */
  refreshIntervalMs?: number
}
```

[packages/context/time-context/src/index.ts:48](#)

@deepseek-ai/dsh-tmux-context

agents · sessionProjections

```
/** Per-turn tmux-location scheduling. Invalid values fail plugin load. */
export interface Config {
  /** Minimum milliseconds between durable injections in one session. Omit or set to 0 to inject on every eligible change. */
  refreshIntervalMs?: number
}
```

[packages/context/tmux-context/src/index.ts:36](#)

@deepseek-ai/dsh-token-meter

sessionProjections

```
/** Token-meter plugin configuration; the fixed estimator has no settings. */
export type TokenMeterConfig = Record<string, never>
```

[packages/llm/token-meter/src/types.ts:13](#)

@deepseek-ai/dsh-tool-bash

tools · shell · systemPrompt · shellEnv

```
/** Configuration for the bash tool. */  
export interface Config {  
  /** Expose `run_in_background` (default true); disabled calls are also rejected. */  
  enableRunInBackground?: boolean  
}
```

[packages/shell/tool-bash/src/index.ts:33](#)

@deepseek-ai/dsh-tool-bash-persistent

tools · terminals

```
/** Configuration for the persistent Bash tool. */  
export interface Config {  
  /** PTY backend used for each owner-isolated persistent shell (default `shell`). */  
  backendType?: string  
  /** Wall-clock limit for one command (default 300000). */  
  timeoutMs?: number  
  /** Maximum returned command-output characters before clipping (default 16000). */  
  maxOutputChars?: number  
  /** Model-facing tool description; deployments may describe their environment. */  
  description?: string  
}
```

[packages/shell/tool-bash-persistent/src/index.ts:432](#)

@deepseek-ai/dsh-tool-fs

tools · fs · systemPrompt

```
/** Plugin config (all optional — `Config` supplies the defaults). */
export interface Config {
  /** Default and maximum number of lines returned by one `read` call. */
  readLimit?: number

  /** Maximum characters returned for a single line before truncation. */
  readMaxLineLength?: number

  /** Maximum bytes returned for the selected lines of one `read` call. */
  readMaxBytes?: number

  /** Files at or above this size stream instead of loading whole into memory. */
  readStreamMinSize?: number
}
```

[packages/fs/tool-fs/src/index.ts:25](#)

@deepseek-ai/dsh-tool-fs-search

tools · [systemPrompt](#) · [subprocess](#)

```
/** Plugin config; over-cap glob sampling is an explicit deployment choice and the remaining fields have defaults. */
export interface Config {
  /** Whether an over-cap `glob` page is sampled across top-level entries instead of taking the modification-time head. */
  sampleOverCapGlobResults: boolean

  /** Max paths one `glob` call retains inline; later paths go to the formatted spill file. */
  globMaxResults?: number

  /** Max flat matches one `grep` call retains inline; later matches go to the formatted spill file. */
  grepMaxMatches?: number

  /** Max bytes retained for one matched-line preview (the cut preserves UTF-8 boundaries). */
  grepMaxLineBytes?: number

  /** Max bytes of one search's serialized `presentationMeta`; trailing groups/paths drop past it so the persisted card stays bounded. */
  searchMetaMaxBytes?: number

  /** Max complete raw `rg` stdout bytes a search will parse; larger raw output fails with `SEARCH_RAW_OUTPUT_OVERFLOW`. */
  rawOutputMaxBytes?: number

  /** Terminate-escalation grace (ms), handed to the subprocess seam and bounded by `MAX_TIMER_DELAY_MS`. */
}
```



```
    graceMs?: number

    /** Max bytes retained for one search's stderr tail; the excerpt is embedded in `SEARCH_*` error messages,
    never shown on success. */
    stderrMaxBytes?: number

    /**
     * Cooperative tool-call timeout budget (ms) on both tools, enforced by
     * `@deepseek-ai/dsh-tool-call-timeout-policy` through `exec.signal`.
     */
    timeoutMs?: number
}
```

[packages/fs/tool-fs-search/src/index.ts:73](#)

@deepseek-ai/dsh-tool-goal

[agents](#) · [goals](#) · [tools](#) · [systemPrompt](#) · [sessionProjections](#)

```
/** Model policy and hard lower bounds for goal-state updates. */
export interface Config {
    /** Minimum admitted goal rounds before the model may self-report `blocked`. */
    blockedAfterConsecutiveRounds?: number
}
```

[packages/goal/tool-goal/src/index.ts:25](#)

@deepseek-ai/dsh-tool-jobs

[tools](#) · [jobs](#) · [systemPrompt](#)

```
/** Configures bounded `job_output` waits and completion-notice delivery. */
export interface Config {
    /** Wait duration applied when `job_output` sets `wait` without `timeout_ms` (default 30s). */
    waitTimeoutMs?: number

    /** Hard cap on any single wait; a larger model-supplied `timeout_ms` is clamped down to it (default 10min). */
    maxWaitTimeoutMs?: number

    /** Whether a completion opens a turn on an idle owner (default `wakeup`). */
    completionDelivery?: CompletionDelivery
}
```

```

/**
 * Turns one owner may have opened by completion wakes before the next
 * notice degrades to injection, reset by any user-authored input (default 3).
 * Bounds the self-exciting chain where a woken turn starts the job whose
 * completion wakes it again.
 */
maxConsecutiveWakes?: number
}

/**
 * How an unreported completion reaches an owner that is already idle: `wakeup`
 * opens a turn for it, `quiet` leaves it pending until something else wakes the
 * owner. A busy owner is injected either way.
 */
export type CompletionDelivery = 'quiet' | 'wakeup'

```

[packages/jobs/tool-jobs/src/index.ts:31](#)

@deepseek-ai/dsh-tool-lsp

tools · lsp · systemPrompt

```

/** Plugin configuration: result caps and the timeout budget. */
export interface Config {
  /** Largest number of rendered locations before an omission marker (default 100). */
  maxLocations?: number
  /** Largest complete rendered result in characters, including truncation metadata (default 16000). */
  maxResultChars?: number
  /** Tool-call timeout budget in ms (default 60000). */
  timeoutMs?: number
}

```

[packages/lsp/tool-lsp/src/index.ts:57](#)

@deepseek-ai/dsh-tool-pwsh

tools · shell · systemPrompt · shellEnv

```
/** Configuration for the pwsh tool. */
export interface Config {
  /** Expose `run_in_background` (default true); disabled calls are also rejected. */
  enableRunInBackground?: boolean
}
```

[packages/shell/tool-pwsh/src/index.ts:51](#)

@deepseek-ai/dsh-tool-pwsh-persistent

tools · terminals

```
/** Configuration for the persistent pwsh tool. */
export interface Config {
  /** PTY backend used for each owner-isolated persistent shell (default `shell`). */
  backendType?: string
  /** Wall-clock limit for one command (default 300000). */
  timeoutMs?: number
  /** Maximum returned command-output characters before clipping (default 16000). */
  maxOutputChars?: number
  /** Model-facing tool description; deployments may describe their environment. */
  description?: string
}
```

[packages/shell/tool-pwsh-persistent/src/index.ts:472](#)

@deepseek-ai/dsh-tool-ralph

tools · workflowEngine · subagents · systemPrompt

```
/** Deployment policy for the fixed Ralph workflow. */
export interface Config {
  /** Fresh structured-output provider used for every round (default `spawn`). */
  subagentProvider?: string
  /** Default and deployment ceiling for one call's round count (default 256). */
  maxRounds?: number
}
```

```
/** Maximum serialized characters in one structured handoff (default 16384). */
maxHandoffChars?: number

/** Maximum characters in a successful parent-facing terminal text (default 16384). */
maxResultChars?: number
}
```

[packages/workflow/tool-ralph/src/index.ts:21](#)

@deepseek-ai/dsh-tool-session-query

tools · systemPrompt · sessionQuery · sessionProjections

```
/** Deployment-owned search count and timeout bounds. */
export interface Config {
  /** Maximum authorized hits returned by one search call. Defaults to 100. */
  maxSearchResults?: number
  /** Cooperative full-text search deadline in milliseconds. Defaults to 30000. */
  searchTimeoutMs?: number
}
```

[packages/session-query/tool-session-query/src/index.ts:28](#)

@deepseek-ai/dsh-tool-skill

agents · tools · skills

```
/** Model-facing skill catalog configuration. */
export interface Config {
  /** Maximum normalized description length rendered in the session catalog; minimum 3. */
  catalogDescriptionMaxLength?: number
}
```

[packages/skill/tool-skill/src/index.ts:61](#)

@deepseek-ai/dsh-tool-str-replace-editor

```
/** Configuration for the string-replacement editor tool. */  
export interface Config {  
  /** Maximum returned view characters before clipping (default 16000). */  
  maxOutputChars?: number  
  /** Model-facing tool description. */  
  description?: string  
}
```

[packages/fs/tool-str-replace-editor/src/index.ts:505](#)

@deepseek-ai/dsh-tool-subagent

```
/** Config: which registered provider this tool delegates to, plus child defaults. */  
export interface Config {  
  /** The `ctx.subagents` provider name to start runs on (e.g. `spawn`, `acp`). */  
  provider: string  
  /**  
   * Model-facing tool name (default `subagent`). Each loaded instance must use  
   * a distinct name.  
   */  
  toolName?: string  
  /**  
   * Sample the Host `subagent-model-selection` user setting for each new  
   * top-level session and inherit that decision in its child sessions.  
   */  
  modelSelectionSettings?: boolean  
  /**  
   * Expose `run_in_background` (default true). Disabled instances omit the  
   * parameter and reject forced background calls.  
   */  
  enableRunInBackground?: boolean  
  /**  
   * Background execution policy (default `one-shot`). `one-shot` defaults calls  
   * to foreground; `continuable` defaults them to background, requires a provider
```

```

* with the `prepareContinuable` capability, and returns the durable child id.
* Follow-up adapters remain independently optional.
*/
backgroundMode?: 'one-shot' | 'continuable'
/**
* Agent options applied to every child; omitted fields use child-loop defaults.
*/
agentOptions?: AgentOptions
/**
* Per-child persona that shadows `deployment:persona`. Requires the
* provider's `persona` capability; omission preserves the deployment persona.
*/
persona?: string
/**
* Tool filter applied to every child. Filtered tools disappear from its
* prompt and reject execution. Requires the provider's `toolFilter`
* capability; unknown names fail startup.
*/
toolFilter?: {
  /** Global tool names the child keeps; everything else is removed. */
  allow?: string[]
  /** Global tool names removed from the child. */
  deny?: string[]
}
/**
* Maximum child depth: a non-negative safe integer (default `3`; `0` forbids
* delegation entirely), or `provider-managed` to send no cap. A numeric cap
* requires the provider's `depthLimit` capability (mount fails loud
* otherwise). The provider checks the calling agent's current depth at every
* start; the tool remains model-visible so runtime policy owns rejection.
* `provider-managed` is for an out-of-process provider whose recursion
* budget belongs to the child runtime or its own deployment.
*/
maxDepth?: number | 'provider-managed'
}

```

[AgentOptions](#)

[packages/subagent/tool-subagent/src/index.ts:48](#)

@deepseek-ai/dsh-tool-terminal

terminals · tools · systemPrompt

```
/** Model-facing terminal tool configuration. */  
export interface Config {  
  /** Expose `run_in_background` and accept background sends (default true). */  
  enableRunInBackground?: boolean  
  /** Maximum UTF-8 bytes in one complete terminal or task-output result. */  
  maxResultBytes?: number  
}
```

[packages/terminal/tool-terminal/src/index.ts:35](#)

@deepseek-ai/dsh-tool-todo

tools

```
/** Model-facing todo tool configuration. */  
export interface Config {  
  /**  
    * Required deployment choice for whether several todos may be `in_progress` at once. True suits  
    * agents that run work concurrently — subagents, background commands, workflow fan-out — and the  
    * description then instructs the model to mark every actively worked task. False restores the  
    * single-active discipline: the description asks for exactly one, and a call marking more is  
    * rejected.  
    */  
  allowParallelInProgress: boolean  
}
```

[packages/todo/tool-todo/src/index.ts:29](#)

@deepseek-ai/dsh-tool-web

tools · web · systemPrompt

```
/** Plugin config: which web tools to register, search bounds, per-tool budgets, and the fetch output cap. */
export interface Config {
  /** Register `web_search`. Defaults to true. */
  search?: boolean
  /** Register `web_fetch`. Defaults to true. */
  fetch?: boolean
  /** Upper bound on sources returned by one `web_search` call. */
  searchMaxResults?: number
  /** Upper bound on queries accepted by one `web_search` call. */
  searchMaxQueries?: number
  /** Cooperative timeout budget (ms) for `web_fetch`. Defaults to 30000. */
  fetchTimeoutMs?: number
  /** Cooperative timeout budget (ms) for `web_search`. Defaults to 30000. */
  searchTimeoutMs?: number
  /** Cap on source characters converted and complete `web_fetch` output characters. Defaults to 200000. */
  fetchMaxOutputChars?: number
}
```

[packages/web/tool-web/src/index.ts:37](#)

@deepseek-ai/dsh-tool-workflow

tools · workflowEngine · systemPrompt

```
/** Config: the model-facing tool name plus result rendering caps. */
export interface Config {
  /** The model-facing tool name to register (default `workflow`). */
  toolName?: string
  /** Rendered-result ceiling, in characters: a longer JSON value is truncated with a notice (default 50000). */
  maxResultChars?: number
}
```

[packages/workflow/tool-workflow/src/index.ts:32](#)

@deepseek-ai/dsh-tools

systemPrompt


```

/** Plugin config: how the registered tools are presented to the model. */
export interface Config {
  /**
   * Model presentation. `native` (default) sends every visible schema; `ptc`
   * sends only `run_code` plus a generated SDK prompt and collapses the
   * executor to the same surface (a model-direct call may only name
   * `run_code`; `run_code` SDK sub-dispatches keep every visible tool); `both`
   * sends both forms. PTC mode requires a `ctx.codeRuntime` whose `language`
   * has a registered SDK renderer (TypeScript or Python) and fail prompt
   * assembly when it is absent or has no renderer. Under `ptc`, native names
   * in `toolOrder` are invalid.
   */
  mode?: ToolPresentationMode
  /**
   * Concurrency cap for a `run_code` program's overlapping sub-calls
   * (default 10, the loop scheduler's own default). Sub-calls follow the
   * native scheduling contract — only calls whose tools classify
   * concurrency-safe overlap; exclusive calls form barriers — so `1`
   * restores strictly serial dispatch. Must be a positive integer.
   */
  maxParallelSubCalls?: number
}

/** How the registry presents its tools to the model (see {@link Config.mode}). */
export type ToolPresentationMode = 'native' | 'ptc' | 'both'

```

[packages/core/tools/src/index.ts:647](https://github.com/deepseek-ai/dsh-typert-loader/blob/main/packages/core/tools/src/index.ts#L647)

@deepseek-ai/dsh-typert-loader

typert · loader

```

/** Additional package artifacts whose owning plugins are nested behind another Loader entry. */
export interface Config {
  /** Exact npm package names that must resolve and export `./typert`. */
  packages?: string[]
}

```

@deepseek-ai/dsh-user-approval

```
/** Plugin config. All optional — `static Config` supplies the defaults. */
export interface Config {
  /**
   * The deployment's default {@link ApprovalPolicy} for sessions without an
   * `approval/policy` override — `ask` delegates to the composed answerers
   * (fail-closed with none); `never` auto-rejects every ask without
   * prompting (the deterministic CI/unattended stance).
   */
  readonly policy?: ApprovalPolicy
}

/**
 * A session's approval policy — what happens to an {@link ApprovalService}
 * ask BEFORE any interactive answerer sees it:
 *
 * - `ask` (the default) — delegate to the composed answerers; with none
 *   composed the chain falls through to the fail-closed `unavailable`.
 * - `never` — never prompt anyone: every ask resolves `rejected`
 *   deterministically. The strict headless stance (CI, unattended runs) and
 *   the policy whose outcome is knowable without asking.
 */
export type ApprovalPolicy = 'ask' | 'never'
```

@deepseek-ai/dsh-web

```
/**
 * Config for the web seam. `searchProvider` / `fetchProvider` pin which provider
 * wins for each capability; both are optional (a single registered usable
 * provider auto-selects). Operational overrides such as environment variables
 * must feed these same fields rather than introduce a hidden priority chain.
```

```

*/
export interface WebRuntimeConfig {
  /** Explicit search provider id. Omitted = auto-select when exactly one usable. */
  readonly searchProvider?: string
  /** Explicit fetch provider id. Omitted = auto-select when exactly one usable. */
  readonly fetchProvider?: string
}

```

[packages/web/web/src/index.ts:55](#)

@deepseek-ai/dsh-web-app

webServer

```

/** Plugin config: composed deployment settings plus per-invocation command-line values. */
export interface Config {
  /** Permit default-browser handoff after the Loader tree settles; an SSH launch suppresses it. */
  openBrowser: boolean
  /** Print the URL line on activation; a non-interactive layer can turn it off. */
  printUrl: boolean
  /**
   * Register the model-visible surface context (the `app:web-surface` prompt
   * section and the `DSH_WEB_URL` bash variable). A one-shot non-interactive
   * layer can turn it off when its user is not in the GUI, so the
   * orientation text would be false.
   */
  surfaceContext: boolean
  /** Explicit `--trusted-host` authorities from this invocation. */
  trustedHosts: string[]
}

```

[packages/bundle/web-app/src/index.ts:44](#)

@deepseek-ai/dsh-web-fetch-http

web

```
/** Plugin config: the provider's transport and size limits plus its `User-Agent` (all defaulted). */
export interface Config {
  /** Maximum response body size in bytes. */
  maxResponseBytes?: number
  /** Maximum decoded body length in characters. */
  maxBodyChars?: number
  /** Default fetch timeout in milliseconds, within Node's timer range. */
  timeoutMs?: number
  /** Maximum number of same-origin redirect hops to follow. */
  maxRedirects?: number
  /** `User-Agent` header sent on every request. */
  userAgent?: string
}
```

[packages/web/web-fetch-http/src/index.ts:32](https://github.com/deepseek-ai/dsh-web-search-deepseek/blob/main/packages/web/web-fetch-http/src/index.ts#L32)

@deepseek-ai/dsh-web-search-deepseek

web

```
/** Plugin config (all optional — `apply` fills env-var and constant defaults). */
export interface Config {
  /** Literal DeepSeek API key; prefer {@link apiKeyEnv} so no secret enters configuration files. */
  apiKey?: string
  /** Credential reference resolved for each search; defaults to `DEEPSEEK_API_KEY`. */
  apiKeyEnv?: string
  /** Anthropic-compatible endpoint base; `/messages` is appended. */
  baseUrl?: string
  /** Anthropic-format model name. Defaults to `deepseek-v4-flash`. */
  model?: string
  /** `anthropic-version` header value. Defaults to `2023-06-01`. */
  apiVersion?: string
  /** Upper bound on generated tokens for the Messages request. Defaults to 4096. */
  maxTokens?: number
  /** Maximum `web_search` server-tool uses per request. Defaults to 5. */
  maxUses?: number
}
```

[packages/web/web-search-deepseek/src/index.ts:46](#)

@deepseek-ai/dsh-web-search-exa

web

```
/** Plugin config (all optional — `apply` fills env-var and constant defaults). */
export interface Config {
  /** Exa API key. Falls back to `EXA_API_KEY`. Empty → provider unavailable. */
  apiKey?: string
  /** Endpoint base; `/search` is appended. Defaults to the public API. */
  baseUrl?: string
  /** Retrieval mode sent as Exa's `type`. Defaults to `auto`. */
  searchType?: 'auto' | 'keyword' | 'neural'
  /** Default result count when a request carries no `maxResults`. Omitted = none. */
  numResults?: number
  /** Highlight sentences requested per result. Defaults to 1. */
  highlightsPerResult?: number
}
```

[packages/web/web-search-exa/src/index.ts:35](#)

@deepseek-ai/dsh-web-search-perplexity

web

```
/** Plugin config (all optional — `apply` fills env-var and constant defaults). */
export interface Config {
  /** Perplexity API key. Falls back to `PERPLEXITY_API_KEY`. Empty → unavailable. */
  apiKey?: string
  /** Endpoint base; `/chat/completions` is appended. Defaults to the public API. */
  baseUrl?: string
  /** Search model name. Defaults to `sonar`. */
  model?: string
  /** Upper bound on generated answer tokens. Defaults to 1024. */
  maxTokens?: number
  /** Recency window sent as `search_recency_filter`. Omitted = no filter. */
}
```

```
searchRecency?: 'day' | 'week' | 'month' | 'year'  
}
```

[packages/web/web-search-perplexity/src/index.ts:30](#)

@deepseek-ai/dsh-webhook-github

webServer · webhookRuntime · credentials

```
/** Required GitHub ingress configuration. */  
export interface Config {  
  /** Adapter instance name carried to rules. */  
  readonly source: string  
  /** Exact absolute route path. */  
  readonly path: string  
  /** Credential reference containing the shared webhook secret. */  
  readonly secretEnv: string  
  /** Positive raw body ceiling in bytes. */  
  readonly maxBodyBytes: number  
}
```

[packages/webhook/webhook-github/src/index.ts:17](#)

@deepseek-ai/dsh-workflow-worker-thread

subagents

```
/** Plugin config (all optional — `static Config` supplies the defaults). */  
export interface Config {  
  /** The `ctx.subagents` provider children run on (default `spawn`). */  
  provider?: string  
  /** Concurrent `agent()` ceiling; `0` (the default) auto-resolves to `min(16, max(1, cores - 2))`. */  
  maxConcurrentAgents?: number  
  /** Total `agent()` calls one run may start — the runaway-loop backstop (default 1000). */  
  maxTotalAgents?: number
```

[packages/workflow/workflow-worker-thread/src/index.ts:32](#)

- @deepseek-ai/dsh-client-ui-layout [packages/client/ui-layout/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-message-feedback [packages/client/ui-message-feedback/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-model-selection [packages/client/ui-model-selection/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-permission-presets [packages/client/ui-permission-presets/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-plan [packages/client/ui-plan/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-reference [packages/client/ui-reference/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-renderer [packages/client/ui-renderer/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-schedule [packages/client/ui-schedule/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-session [packages/client/ui-session/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-settings [packages/client/ui-settings/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-settings-general [packages/client/ui-settings-general/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-settings-models [packages/client/ui-settings-models/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-settings-plugin-inventory [packages/client/ui-settings-plugin-inventory/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-settings-plugins [packages/client/ui-settings-plugins/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-sidebar [packages/client/ui-sidebar/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-skill [packages/client/ui-skill/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-subagent [packages/client/ui-subagent/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-theme [packages/client/ui-theme/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-tool [packages/client/ui-tool/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-trajectory [packages/client/ui-trajectory/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-user-questions [packages/client/ui-user-questions/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-workflow-run [packages/client/ui-workflow-run/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-workspace [packages/client/ui-workspace/src/index.ts](#) 
- @deepseek-ai/dsh-command-compact —  `commands` · `compact` [packages/compaction/command-compact/src/index.ts](#) 
- @deepseek-ai/dsh-command-feedback —  `commands` [packages/feedback/command-feedback/src/index.ts](#) 
- @deepseek-ai/dsh-command-goal —  `commands` · `goals` [packages/goal/command-goal/src/index.ts](#) 
- @deepseek-ai/dsh-commands [packages/interaction/commands/src/index.ts](#) 
- @deepseek-ai/dsh-cordis-client-runner [packages/extensions/cordis-client-runner/src/index.ts](#) 
- @deepseek-ai/dsh-deepseek-llm-api-extensions [packages/llm/deepseek-llm-api-extensions/src/index.ts](#) 
- @deepseek-ai/dsh-experimental-client-ui-agent-team [packages/experimental/client-ui-agent-team/src/index.ts](#) 
- @deepseek-ai/dsh-fs-e2b —  `e2b` [packages/e2b/fs-e2b/src/index.ts](#) 
- @deepseek-ai/dsh-fs-observation-policy [packages/fs/fs-observation-policy/src/index.ts](#) 
- @deepseek-ai/dsh-goal-round-driver —  `agents` · `goals` · `sessions` [packages/goal/goal-round-driver/src/index.ts](#) 
- @deepseek-ai/dsh-host-directory-picker-auto —  `webServer` · `loader` [packages/host/directory-picker-auto/src/index.ts](#) 
- @deepseek-ai/dsh-host-directory-picker-native [packages/host/directory-picker-native/src/index.ts](#) 
- @deepseek-ai/dsh-host-plugin-inventory —  `loader` [packages/host/plugin-inventory/src/index.ts](#) 
- @deepseek-ai/dsh-llm [packages/llm/llm/src/index.ts](#) 
- @deepseek-ai/dsh-lsp [packages/lsp/lsp/src/index.ts](#) 

- @deepseek-ai/dsh-session-query — [SessionQueryEngine [packages/session-query/session-query/src/index.ts](#)]
- @deepseek-ai/dsh-settings — [Settings [packages/settings/settings/src/index.ts](#)]
- @deepseek-ai/dsh-shell — [ShellExecutor [packages/shell/shell/src/index.ts](#)]
- @deepseek-ai/dsh-spill — [SpillStore [packages/spill/spill/src/index.ts](#)]
- @deepseek-ai/dsh-subprocess — [SubprocessRuntime [packages/subprocess/subprocess/src/index.ts](#)]
- @deepseek-ai/dsh-workflow — [WorkflowEngine [packages/workflow/workflow/src/index.ts](#)]

cordis.yml |

--	--	--	--	--	--

- @deepseek-ai/dsh-agent-loop-testkit [packages/test-support/agent-loop-testkit/src/index.ts](#) 
- @deepseek-ai/dsh-anonymous-user-id [packages/identity/anonymous-user-id/src/index.ts](#) 
- @deepseek-ai/dsh-app-boot [packages/boot/app-boot/src/index.ts](#) 
- @deepseek-ai/dsh-atomic-write [packages/util/atomic-write/src/index.ts](#) 
- @deepseek-ai/dsh-base [packages/bundle/base/src/index.ts](#) 
- @deepseek-ai/dsh-brand [packages/util/brand/src/index.ts](#) 
- @deepseek-ai/dsh-client-store [packages/client/store/src/index.ts](#) 
- @deepseek-ai/dsh-client-test-runtime [packages/test-support/client-runtime/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-primitives [packages/client/ui-primitives/src/index.ts](#) 
- @deepseek-ai/dsh-client-ui-slots [packages/client/ui-slots/src/index.ts](#) 
- @deepseek-ai/dsh-client-web [packages/client/web/src/index.ts](#) 
- @deepseek-ai/dsh-cmdline [packages/boot/cmdline/src/index.ts](#) 
- @deepseek-ai/dsh-deque [packages/util/deque/src/index.ts](#) 
- @deepseek-ai/dsh-experimental-agent-team-profile [packages/experimental/agent-team-profile/src/index.ts](#) 
- @deepseek-ai/dsh-experimental-agent-team-web-profile [packages/experimental/agent-team-web-profile/src/index.ts](#) 
- @deepseek-ai/dsh-experimental-webworker-packer [packages/experimental/webworker-packer/src/index.ts](#) 
- @deepseek-ai/dsh-experimental-webworker-runtime [packages/experimental/webworker-runtime/src/index.ts](#) 
- @deepseek-ai/dsh-home-paths [packages/util/home-paths/src/index.ts](#) 
- @deepseek-ai/dsh-hook-protocol [packages/hooks/hook-protocol/src/index.ts](#) 
- @deepseek-ai/dsh-launch-environment [packages/util/launch-environment/src/index.ts](#) 
- @deepseek-ai/dsh-llm-mock-server [packages/test-support/llm-mock-server/src/index.ts](#) 
- @deepseek-ai/dsh-loader-smoke [packages/test-support/loader-smoke/src/index.ts](#) 
- @deepseek-ai/dsh-native-command [packages/util/native-command/src/index.ts](#) 
- @deepseek-ai/dsh-output-retention [packages/util/output-retention/src/index.ts](#) 
- @deepseek-ai/dsh-sandbox-windows-acl [packages/sandbox/sandbox-windows-acl/src/index.ts](#) 
- @deepseek-ai/dsh-scope [packages/core/scope/src/index.ts](#) 
- @deepseek-ai/dsh-sdk-client [packages/sdk/client/src/index.ts](#) 
- @deepseek-ai/dsh-sdk-minimal [packages/bundle/sdk-minimal/src/index.ts](#) 

- @deepseek-ai/dsh-sdk-protocol [packages/sdk/protocol/src/index.ts](#) 
- @deepseek-ai/dsh-session-snapshot [packages/test-support/session-snapshot/src/index.ts](#) 
- @deepseek-ai/dsh-session-telemetry [packages/session/session-telemetry/src/index.ts](#) 
- @deepseek-ai/dsh-session-title-llm [packages/session/session-title-llm/src/index.ts](#) 
- @deepseek-ai/dsh-subagent-in-process-driver [packages/subagent/subagent-in-process-driver/src/index.ts](#) 
- @deepseek-ai/dsh-timeout [packages/util/timeout/src/index.ts](#) 
- @deepseek-ai/dsh-typer-generator [packages/typer/generator/src/index.ts](#) 
- @deepseek-ai/dsh-typer-protocol [packages/typer/protocol/src/index.ts](#) 
- @deepseek-ai/dsh-typer-registry [packages/typer/registry/src/index.ts](#) 
- @deepseek-ai/dsh-util-crypto [packages/util/crypto/src/index.ts](#) 
- @deepseek-ai/dsh-util-time [packages/util/time/src/index.ts](#) 
- @deepseek-ai/dsh-util-values [packages/util/values/src/index.ts](#) 
- @deepseek-ai/dsh-util-workspace-path [packages/util/workspace-path/src/index.ts](#) 
- @deepseek-ai/dsh-win32-process [packages/subprocess/win32-process/src/index.ts](#) 

 #1

  3  2026 17:09:57

  3  2026 17:36:27