MCP SERVER | NO CODE | CLOUD HOSTED

D&D Character Stat Roller MCP for AI Agents

Statistically Comparing Dice Rolling Mechanics and Attribute Generation

D&D Character Stat Roller is a statistical engine for RPG designers and players. It runs thousands of randomized dice rolls to compare how different character attribute generation methods—like point buy versus rolling 4d6—actually perform. You get hard data on consistency, average scores, and the probability of building powerful characters.

A+ Quality Score 100/100

dnd | ttrpg | dice-rolling | probability | character-creation

MAINTAINED BY VINKIUS CLOUD FOR AI AGENTS · JUNE 2026
vinkius.com/mcp/dd-character-stat-roller

SCAN TO VIEW

The connectivity layer between AI and the world's software.

Vinkius sits between AI and every application. All communication passes through Vinkius Cloud via the Model Context Protocol (MCP) — with governance, observability, and security at every layer.

Your AI Connections Run Through Vinkius Cloud

The world's largest managed MCP catalog

Vinkius is the connectivity layer where AI connects to the software your business already runs. We handle the hosting, the security, the credentials, the uptime — you get agents that actually do things.

We operate the world's largest managed MCP catalog. Major SaaS platforms, CRMs, databases, and cloud providers — running, monitored, production-ready. This MCP server is hosted and maintained by the Vinkius Cloud for AI Agents.

> *The agent doesn't manage credentials, doesn't manage uptime, doesn't manage security. Vinkius does.*
>
> — Architecture principle

Four Pillars of the Vinkius Runtime

01 — Security by design

Credentials stay encrypted at rest via AES-256. The AI agent never touches raw keys — they're injected into a sandboxed V8 isolate at runtime. Actions are logged, and connections have an emergency kill switch.

02 — Built on MCP Fusion

This MCP server was built with **MCP Fusion**, the open-source framework (Apache 2.0) that powers the entire Vinkius catalog. Schema-as-firewall strips undeclared fields, compiled PII redaction runs at zero overhead, and cryptographic lockfiles produce git-diffable audit trails.

03 — Deterministic observability

Eight immutable metrics per endpoint: request volume, p95 latency, error rate, active connections, cost attribution. A live payload feed logs every tool call with mutation detection.

04 — Autonomous operations

Servers are deployed, monitored, and patched autonomously. New capabilities and security patches ship weekly. Zero-downtime deployments ensure continuous availability across all managed MCP servers.

AES-256	Ed25519	24h TTL	V8 Isolate
Encryption at rest	PKI vault signatures	Ephemeral session keys	Sandboxed execution

One Token. Instant Access.

Every MCP server on Vinkius is accessed through a **Connection Token**. Tokens are generated in the cloud dashboard and produce a unique MCP endpoint URL. Paste this URL into any MCP-compatible client — no SDK required.

A single token can serve **multiple AI clients simultaneously**, or you can issue separate tokens per client for granular access control. Each token tracks its own request count, last activity timestamp, and can be individually enabled or revoked.

MCP ENDPOINT

```
https://edge.vinkius.com/{token}/mcp
```

Claude · Cursor · VS Code · Windsurf · Grok · Gemini

Security Is the Architecture

Security in Vinkius is not a feature — it's the foundation of the runtime. The gateway enforces multiple independent protection layers between AI agents and third-party APIs.

01 — Ed25519 PKI Vault

Every workspace has an Ed25519 Master Key. Session keys are generated ephemerally (24h TTL) and signed by the Master Key. Credentials never leave the vault boundary.

02 — V8 Isolate Sandboxing

Tool code runs inside isolated-vm V8 isolates with 64 MB memory caps and per-request timeouts. No filesystem access, no network access except through the SSRF-guarded fetch bridge.

03 — SSRF Guard

All outbound HTTP requests are DNS-resolved and validated before execution. Private IP ranges (10.x, 172.16-31.x, 192.168.x, AWS metadata 169.254.x) are blocked at the network layer.

04 — DLP & PII Redaction

A ResponseGuard pipeline intercepts every tool response. Configurable redaction patterns strip sensitive fields (emails, SSNs, card numbers) before data reaches the AI agent.

05 — Cryptographic Audit Trail

Every request is signed into a SHA-256 hash chain with Ed25519 signatures. Events form a tamper-proof, SIEM-exportable forensic record.

06 — Honeytoken Trap System

Phantom credentials are injected into isolated environments. If a honeytoken is used outside Vinkius infrastructure, the server is quarantined instantly.

Emergency Kill Switch

EU AI Act Art. 14(1) Compliant

The kill switch is an **emergency halt** mechanism — not a simple toggle. When triggered, it executes three actions atomically:

01 — Server deactivated

The MCP server is immediately taken offline across the entire cluster.

02 — All tokens revoked

Every connection token is invalidated. Total lockout — reconnection blocked until new tokens are issued.

03 — WebSocket connections killed

Active connections terminated via Redis pubsub broadcast. Propagates to every runtime node in the cluster.

Full Visibility. Zero Guesswork.

The Vinkius cloud dashboard includes a full MCP Governance suite — real-time analytics and security controls for production AI operations.

Control Plane

KPI dashboard with request volume, latency, success rate, token consumption, and AI-generated operational briefings.

Agent Activity

Which AI clients are connecting, how often, and what they're doing — real-time session tracking.

FinOps

Cost tracking per tool, payload compression savings, budget optimization signals, and consumption trends.

Tool Health

Slowest and most error-prone tools, with actionable root-cause insights and performance baselines.

Firewall & DLP

PII redaction activity, sensitive data protection counters, and security event timeline.

Incident Log

Error trends, failure rates, status-code breakdowns, and forensic audit trail access.

> Get started at **cloud.vinkius.com** — connect your AI agent in under 60 seconds.

D&D Character Stat Roller MCP

4 tools available
Cloud-hosted on Vinkius

Need to know if your game's character creation rules are balanced? This MCP runs statistical simulations comparing various D&D attribute generation methods. Instead of guessing, you simulate thousands of rolls for mechanics like `heroic_4d6`, `classic_3d6`, and standard point-buy systems. Your AI client executes these large-scale randomized tests, letting your agent analyze the math behind character power. You can calculate average scores to see which system provides the most consistent results, or determine the exact probability of a player creating a high-strength character. The Vinkius catalog makes this powerful statistical tool accessible alongside thousands of others, so you don't have to stitch together multiple services just for dice rolling math.

Core Capabilities

01 — Simulate Character Builds

Runs massive randomized tests on a specific character generation method to gather raw data.

02 — Calculate Attribute Averages

Determines the mean score for any attribute across a large set of simulation results.

03 — Compare System Power

Generates detailed reports comparing two different character creation datasets side-by-side.

04 — Assess High-End Potential

Figures out the probability of a generated character exceeding a specific power threshold.

One Click on Vinkius — From Prompt to Execution

Available at **vinkius.com/mcp/dd-character-stat-roller** — connect your AI agent in three steps.

01 You tell your AI client which generation method you want to test (e.g., 'point buy' or '4d6').

02 This MCP runs 10,000 randomized iterations for that method, generating a massive dataset of potential character attributes.

03 Your agent then uses the data to calculate averages or compares it against another system to give you an actionable report.

The bottom line is: You feed it your rules, and it gives you the hard statistical proof of whether those rules hold up under pressure.

Built For

TTRPG designers need this to prove their mechanics aren't broken. Advanced players use it when they want to know if a new system truly rewards careful planning or just luck. If your game needs math, not vibes, you're in the right place.

TTRPG Designer

Uses this MCP to compare mechanics like standard arrays versus rolling rules to ensure character creation is balanced across all six attributes.

Game Master (GM)

Checks if the rules for a new campaign setting are mathematically sound before running it in a session, preventing unexpected power gaps.

System Developer

Validates dice rolling mechanics by simulating thousands of data points to confirm statistical consistency and reliability.

What Changes When You Connect

01 Stop relying on gut feeling. Use the `simulate_method` tool to run 10,000 rolls per mechanic, giving you hard data instead of assumptions about character strength.

02 Pinpoint weak points immediately. The MCP lets your agent use `evaluate_power_threshold` to calculate the exact chance that a player can build an overpowered character in a given system.

03 Compare mechanics objectively. Instead of debating rules, let the numbers talk. Use `generate_method_comparison_report` to see exactly where 'point buy' outperforms 'standard array', and vice versa.

04 Understand core consistency. If you need to know what average strength is mathematically possible in a system, use `analyze_attribute_averages` to get precise mean scores for all six attributes.

05 Save days of playtesting time. You can validate your entire character creation process in minutes, using this MCP's statistical power.

Real-World Applications

Is the 4d6 method actually balanced?

A designer needs to know if their new 'heroic 4d6' rolling rule is consistent. They ask their agent to run a simulation using `simulate_method` and then use `analyze_attribute_averages` on the results. The report proves that while exciting, it generates significantly lower average scores than the old system.

Comparing Point Buy vs. Standard Array

A group is deciding between two established systems. They ask their agent to run a comparison using `generate_method_comparison_report`. The resulting report highlights that while 'point buy' gives higher peak scores, 'standard array' offers much more stable baseline consistency across all attributes.

Checking for Overpowered Builds

A player is worried the game allows too many powerful characters. They ask their agent to use `evaluate_power_threshold` on a specific system, determining that the probability of an 'outperforming' build is actually lower than they thought.

Validating New Dice Mechanics

A developer introduces a complex new rolling rule. They use `simulate_method` to generate 10,000 data points and then ask the agent to calculate the mean using `analyze_attribute_averages`, confirming if the average score matches their initial mathematical predictions.

Patterns to Avoid

Treating dice rolls as random chance

✗ AVOID

Assuming that because a die roll *can* be high, it means most rolls will be high. This ignores the statistical distribution and true average.

✓ INSTEAD

Run simulations using `simulate_method` to gather thousands of data points. Then, use `analyze_attribute_averages` to find the reliable mean score for that method.

Comparing systems manually

✗ AVOID

Trying to compare two character creation methods by rolling them out a dozen times in a spreadsheet and hoping for the best. This sample size is meaningless.

✓ INSTEAD

Use `generate_method_comparison_report` to run both mechanical comparisons across 10,000+ iterations, giving you statistically valid data.

Focusing only on single high rolls

✗ AVOID

Being distracted by the few highest possible scores and ignoring what the typical character will actually look like. This misses core consistency.

✓ INSTEAD

Use `evaluate_power_threshold` to set a specific power metric (e.g., Strength 15) and calculate the precise probability of achieving it, giving you a stable benchmark.

The Right Fit

You should use this MCP if your goal is mathematical certainty about character balance. Use it when you need to compare how different rule sets perform under identical statistical pressure. For example, comparing the mean score of 'point buy' versus 'standard array' requires these tools. Don't use this if you just want quick inspiration or narrative flair; these are for mechanics validation only. If your problem is about *how* characters interact (e.g., combat flow), that's a different kind of modeling. But if the issue is strictly attribute generation, and you need to prove what's mathematically possible—use this MCP.

D&D Character Stat Roller: Pinpointing Attribute Balance in TTRPG Design

Right now, designing character stats often involves a lot of guesswork. You run a few rolls with your friends and say, 'Yeah, that feels right.' Then you copy-paste those results into a spreadsheet to calculate an average, which is obviously wildly inaccurate because the sample size is too small.

With this MCP, your agent runs 10,000 randomized simulations for any system you throw at it. You get a reliable dataset and can immediately use tools like `analyze_attribute_averages` to find the true mean score. The outcome isn't 'it feels right,' it's proven.

D&D Character Stat Roller: Measuring Power Consistency in Character Creation

The biggest time suck is repeatedly checking if a new mechanic—say, swapping 4d6 for a modified roll—actually changes the character's power level significantly. You're stuck running small test groups and hoping nothing major breaks.

This MCP lets you bypass that tedious cycle. By using `generate_method_comparison_report` or checking against specific benchmarks with `evaluate_power_threshold`, you get a definitive, statistically robust answer in minutes. It tells you exactly what's different.

D&D Character Stat Roller: 4 Tools for Attribute Generation Statistics

Use these tools to simulate, analyze averages, compare entire systems, and determine the probability of specific high-power builds.

#	TOOL	DESCRIPTION
01	`simulate_method`	Runs a massive randomized test for any character attribute generation rule you specify.
02	`analyze_attribute_averages`	Calculates the average score for specific attributes using your existing simulation data set.
03	`generate_method_comparison_report`	Pulls together a summary report showing how two different character creation systems stack up against each other.
04	`evaluate_power_threshold`	Determines the specific probability of creating a character that meets or exceeds a high-end power target.

See It in Action

Real prompts you can use once this MCP is connected to your AI agent through Vinkius Cloud.

U Compare the average strength and constitution scores for point buy vs standard array.

Comparison Report: Primary Attributes

ATTRIBUTE	POINT BUY AVERAGE	STANDARD ARRAY AVERAGE
Strength	12.95	11.78
Constitution	13.01	12.45

The data shows that while the standard array is consistent, point buy leads to a significantly higher average score in both primary physical attributes.

U What's the chance of getting at least two stats above 16 with 3d6?

Probability Analysis: Classic 3d6

Based on 10,000 randomized rolls, the probability of having 2 or more attributes rated 16+ is **9.5%**. This means that for every 100 characters you build using this method, about 9 to 10 will hit that high threshold.

U Simulate a few rolls using the heroic 4d6 rule and find the average.

Simulation Complete: Heroic 4d6

- Total Iterations Run: 10,000
- Calculated Mean Strength: **12.81**
- Standard Deviation: 3.11

The average strength score is 12.81. This figure represents the true mathematical center of the distribution for this method.

Frequently Asked Questions

01 **How does D&D Character Stat Roller help me balance my game rules?**

It provides statistical proof that your mechanics are balanced. Instead of guessing, you run simulations to compare different methods—like point buy vs. rolling 4d6—and see which one generates the most consistent and reliable results across all attributes.

02 **Can this MCP tell me if my character class is too strong?**

Yes, it helps you gauge potential power gaps. You can run simulations to determine the exact probability of a character achieving a high-end score on key metrics, helping you balance overpowered build possibilities.

03 **What kind of data does D&D Character Stat Roller provide?**

It delivers hard numbers: average scores, standard deviations, and comparative reports. You get to see the mathematical consistency of your rules across thousands of simulated characters, not just a handful.

04 **Do I need to be a designer to use D&D Character Stat Roller?**

No. Any player who wants to compare dice rolling systems or validate custom homebrew mechanics can use it. It's perfect for anyone serious about the math behind character creation.

05 **Is this better than just using an online dice roller?**

A standard dice roller only gives you random outcomes, which is fun but useless for balance checking. This MCP runs thousands of rolls to find the *average* and *probability*, giving you deep statistical insight.

Go Live in 60 Seconds

Get your connection token from **cloud.vinkius.com**, then paste the endpoint URL into any MCP-compatible client.

YOUR MCP ENDPOINT

```
https://edge.vinkius.com/[TOKEN]/mcp
```

CLIENT	WHERE TO CONFIGURE
Claude AI	Profile → Customize → Connectors → "+" → Add custom connector → Paste endpoint
Cursor	Settings → Features → MCP Servers → "+ Add New MCP Server" → Type: SSE → Paste endpoint
VS Code	Ctrl/Cmd+Shift+P → "MCP: Add Server" → add `"dd-character-stat-roller": { "url": "…" }`
Windsurf	MCP Settings → `mcp_settings.json` → Add endpoint URL
ChatGPT	Settings → Tools & plugins → Add MCP server → Paste endpoint
Gemini	Extensions → Add MCP Server → Paste endpoint URL

ASK AN AI ABOUT THIS

Let your preferred AI explain this MCP server

- Ask ChatGPT
- Ask Claude
- Ask Perplexity
- Ask Gemini
- Ask Grok

READY TO CONNECT

D&D Character Stat Roller is live on Vinkius Cloud.

Get your connection token, paste it into your AI agent, and start building. No SDK. No deployment. Just results.

Start at cloud.vinkius.com →

vinkius.com · support@vinkius.com

INDEPENDENT PLATFORM DISCLAIMER

Vinkius is an independent platform and is not affiliated with, endorsed by, sponsored by, verified by, or otherwise authorized by D&D Character Stat Roller. All third-party trademarks, logos, and brand names are the property of their respective owners. Their use in this document is strictly for informational purposes to identify service compatibility and interoperability.

DOCUMENT INFORMATION

Generated	June 2026
MCP Server	D&D Character Stat Roller MCP
Server ID	`019f0f8f-ed8f-73a9-be27-0c634856b3dd`
Platform	Vinkius Cloud for AI Agents
Endpoint	`https://edge.vinkius.com/{token}/mcp`

LICENSE & USAGE

This document is generated automatically by the Vinkius PDF Engine. Content reflects the MCP server configuration at the time of generation and may change as updates are deployed. For the most current information, visit vinkius.com/mcp/dd-character-stat-roller.