CDATA

Whitepaper

Enterprise Security Best Practices Guide

Building Secure MCP Architectures for AI Data Connectivity

March 2026

What This Guide Covers

- Understanding MCP security risks and why managed platforms matter
- Security architecture fundamentals for enterprise MCP deployments
- Identity-first access controls and RBAC enforced at the source
- Governance, audit trails, and compliance best practices
- Answering common security review questions

For IT Leaders, Security Teams, and Data Architects

Executive Overview

As organizations deploy AI agents and copilots to automate business processes, three critical challenges emerge: connectivity to enterprise data sources, context about what that data means, and control over what AI can access and do.

CData Connect AI is the first managed Model Context Protocol (MCP) platform designed to address all three requirements with enterprise-grade security built in from the ground up. This guide provides IT leaders and security teams with actionable best practices for building secure MCP architectures that scale.

Connect AI Security Foundation

CData has completed third-party SOC 2 Type II and ISO/IEC 27001:2022 compliance audits. Connect AI provides comprehensive security including OAuth 2.1 with PKCE, SSO integration, AES-256 encryption at rest, and TLS 1.3 in transit. Data is never stored or copied—all queries execute in real-time against source systems, maintaining data sovereignty and compliance.

Understanding MCP Security Risks

The Model Context Protocol enables AI agents to interact with enterprise systems in real-time. While this unlocks significant value, recent industry research has identified security concerns that organizations must address before production deployment.

Why Security Matters in AI Data Connectivity

MCP introduces new security considerations that traditional data integration approaches don't fully address:

- AI agents behave non-deterministically, complicating traditional authentication workflows
- Prompt injection risks can potentially trick agents into exposing sensitive data
- Token and credential management across multiple systems requires careful orchestration
- MCP protocol itself only has basic optional security controls at the server level; the way MCP is implemented is more important to provide security controls across AI users

Risks of Unmanaged MCP Deployments

Organizations building their own MCP infrastructure or using community servers face significant security challenges:

Supply Chain Risks

Installing local MCP servers is equivalent to running arbitrary code from often-unvetted sources. Current MCP server distribution relies heavily on unofficial registries, with installation patterns that resemble the "pipe curl to bash" anti-pattern—no pinning, signing, or package locking.

Registry and Trust Issues

Public MCP server registries face familiar open-source ecosystem risks:

- Typosquatting: Malicious packages with names similar to popular ones
- Impersonation: Developers falsely claiming to represent known projects
- Rug pulls: Packages updated with malicious code after gaining adoption
- Account takeovers: Compromised developer accounts pushing malicious updates

Current registries offer weak or inconsistent trust signals. "Official" and "verified" labels often do not confirm developer identity or establish trusted connections between servers and the products they claim to represent.

Remote Server Vulnerabilities

While remote MCP servers don't run locally, they can still enable remote code execution, credential theft, or unauthorized access through other tools available to the client. Vendor risk is also a factor—remote servers may store or process sensitive authentication data and customer information.

Client-Side Risks

MCP clients vary significantly in security posture. Many allow auto-running tools without human approval, which implicitly trusts tool responses and increases the blast radius of compromised servers. Additional risks include tool name conflicts, slash command hijacking, and indirect prompt injection from untrusted content.

How Connect AI Addresses These Risks

Risk Category	Unmanaged MCP	Connect AI Solution
Prompt Injection	Permission creep and default read/write scope	Unauthorized prompt actions are handled with custom MCP tools to restrict what actions a specific agent can take
Supply Chain	Unvetted code from public code repositories	Managed, audited platform with SOC 2 Type II certification and third-party penetration testing
Authentication	Custom authentication at the individual server level	Centralized OAuth 2.1 with PKCE, SSO integration across MCP servers and AI users
Access Control	Rebuild permissions in each system	RBAC passthrough honors existing source controls
Data Security	Data copies across multiple systems	In-place access—data is never replicated from source system
Audit Logging	Inconsistent or missing logs	Comprehensive trails: user, query, timestamp, errors
Credential Management	Scattered tokens and secrets	Centralized, encrypted credential storage
Governance	No standardized controls	CRUD scoping, workspaces, derived views, custom MCP tools, and toolkits

Industry Recommendation: Use Managed MCP Platforms

Security researchers recommend centralizing MCP server usage through managed platforms that provide audit logging, monitoring, guardrails, and governance controls. This approach creates a single point of control for security oversight rather than distributed, unmanaged servers.

Security Architecture Fundamentals

Connect AI's security architecture is built on three core principles: identity-first access, in-place data connectivity, and comprehensive audit logging.

In-Place Data Access

Unlike traditional integration patterns that extract and replicate data, Connect AI accesses data in place entirely copying from source systems. CData never stores data in persistent storage or caching, only temporarily transmitting data at time of query. This provides several security advantages:

- Source system access controls remain authoritative and unchanged
- No data copies to secure, manage, or synchronize
- Data sovereignty requirements are inherently satisfied
- Governance controls at the source automatically apply to AI access

For organizations with on-premises data, the Connect Gateway extends these same guarantees behind the firewall. The gateway runs as a lightweight agent within your infrastructure and establishes a secure, outbound-only connection to Connect AI. On-premises systems are never directly reachable from the public internet, and the same audit trails, RBAC passthrough, and permission controls that apply to cloud sources apply equally to on-premises systems connected through the gateway.

Identity-First Security Model

Connect AI implements a passthrough authentication model where source system RBAC flows through unchanged. When a user or AI agent requests data, the platform:

- Authenticates the request using the configured identity provider
- Passes credentials through to the source system
- Executes the query with that user's individual permissions or inherited from role-level permissions
- Logs the action under the authenticated identity for audit purposes

This approach ensures users and AI agents only access data they're authorized to see, without requiring IT to rebuild permission structures.

Platform Certifications and Compliance

Certification	Description
SOC 2 Type II	Independent audit validates operational effectiveness of security controls over time. Covers all CData products including Connect AI.
ISO/IEC 27001:2022	International standard for information security management systems (ISMS), demonstrating systematic approach to managing sensitive data.
GDPR Compliant	Platform design supports GDPR requirements through data minimization, in-place access, and comprehensive audit capabilities.

Security Testing

CData engages independent third-party penetration testers to conduct regular security assessments of Connect AI. Additionally, our internal security team performs ongoing penetration testing and vulnerability assessments between external audits. Summary findings from third-party assessments are available upon request under NDA.

Authentication Best Practices

Proper authentication configuration is the foundation of a secure MCP deployment. Connect AI supports multiple authentication methods to integrate with your existing identity infrastructure.

Single Sign-On (SSO) Integration

Connect AI supports Single Sign-On through industry-standard providers. When SSO is enabled, users authenticate through your chosen provider instead of Connect AI login credentials.

Supported SSO Providers:

- SAML 2.0 and OpenID Connect
- Microsoft Entra ID (formerly Azure AD)
- Google Workspace
- Active Directory Federation Services (ADFS) and AD/LDAP
- Ping Federate and Okta Workforce Identity Cloud

Recommendation: Enable SSO for Production Deployments

For enterprise deployments, we recommend enabling SSO to centralize authentication, enforce existing security policies, simplify user provisioning, and maintain consistent audit trails across your identity infrastructure. Contact your CData account team to enable SSO for your account.

OAuth 2.1 with PKCE

Connect AI implements OAuth 2.1 with Proof Key for Code Exchange (PKCE) to protect authorization flows. This modern authentication standard provides enhanced security for MCP connections by:

- Preventing authorization code interception attacks
- Eliminating the need for client secrets in public clients
- Supporting secure token refresh without exposing credentials

User Credentials Mode

For data sources that support it, Connect AI offers User Credentials mode, which forces each user to authenticate with their own credentials rather than shared connection credentials.

Supported Data Sources: Salesforce, Snowflake, SharePoint, Workday, Sage Intacct, RedShift, and Paylocity. Contact your account manager to turn on User Credentials mode for any other source.

When enabled, User Credentials mode counts as a single connection slot regardless of how many users authenticate, making it both secure and cost-effective.

SCIM 2.0 Support: Automated Identity Lifecycle Management

Connect AI supports SCIM 2.0 provisioning, enabling organizations to synchronize users, groups, and permissions directly from identity providers including Okta, Microsoft Entra ID, and Ping Identity. This provides benefits such as:

- Automated user lifecycle management eliminates manual provisioning and deprovisioning workflows
- Orphaned credentials risk is eliminated; access is revoked in Connect AI the moment a user is removed from the identity provider
- Aligns with mandatory automated access control requirements under SOC 2 and ISO 27001, removing a common blocker for security team sign-off

Access Control Configuration

Connect AI provides multiple layers of access control to implement least-privilege principles and ensure users and AI agents only access the data they need.

Identity Passthrough

The identity passthrough model enforces role-based access control (RBAC) at the source. Rather than rebuilding permissions in a new system, Connect AI honors existing access controls from your source systems.

Benefits of RBAC Passthrough:

- No need to duplicate permission structures across systems
- Source system permission changes automatically apply to AI access
- Reduces administrative overhead and configuration drift
- Simplifies compliance auditing with single source of truth for permissions

CRUD Operation Scoping

Beyond source system permissions, Connect AI allows administrators to further restrict what operations AI agents can perform. Support for create, read, update, and delete operations can be individually controlled per user or per connection.

Operation	Use Case	Recommendation
Create	New record creation, data entry automation	Enable selectively for automation workflows
Read	Data retrieval, reporting, analytics	Enable by default for most use cases
Update	Record modification, status updates	Restrict to specific trusted workflows
Delete	Record removal, data cleanup	Disable unless explicitly required

User Permission Configuration

Connect AI includes a management layer with granular controls over data types available to individual users:

- Connection-level access: Control which data sources each user can access
- Schema-level restrictions: Limit access to specific schemas within a connection
- Operation scoping: Control create/read/update/delete permissions independently
- Role-based assignment: Group users with similar access requirements

Organizing Data for Secure Access

Connect AI provides features to organize and curate data access, enabling you to expose only the data AI agents need while maintaining security boundaries.

Workspaces

Workspaces create a data catalog from connected sources, bundling related data items in a scalable way. This reduces data silos while promoting controlled access across teams.

Security benefits:

- Bundle tables, views, and derived views into logical collections
- Expose data without transformations, preserving original column names and properties
- Reference items by unique aliases for clarity
- Full CRUD capabilities available for workspace tables when appropriate

Best Practice: Create Purpose-Specific Workspaces

Create predefined, multi-source data collections with custom tools tailored to specific use cases. This approach combines targeted datasets with the universal toolset to boost both performance and security for specialized agent deployments. For example, create a "Sales Analytics" workspace that includes only the Salesforce objects and Snowflake tables needed for sales reporting.

Derived views

Derived views are saved queries that dynamically populate data when accessed. They provide a powerful mechanism for controlling what data AI agents can see by pre-filtering sensitive information.

Security applications:

- Pre-filter data to exclude sensitive columns or records
- Create consistent, curated data subsets for specific use cases
- Reduce attack surface by limiting exposed data
- Enforce data masking or aggregation requirements

Permission Inheritance: Saved derived views appear in the derived views list for all users in the account. However, users who query derived views must have appropriate permissions for each data source that the derived view queries. If a user lacks appropriate permissions, they receive an error—ensuring derived views don't bypass source system access controls.

Custom Tools and Toolkits

Connect AI provides a three-tier tool architecture that gives organizations precise control over what agents can see and do. Each tier serves a distinct purpose in the governance stack:

- Universal Tools provide a compact, schema-aware interface that works consistently across all 350+ connected systems. Instead of exposing hundreds of system-specific operations, agents receive a normalized set of tools that reduce token consumption and limit unnecessary data exposure during exploratory queries.
- Custom Tools allow organizations to define purpose-built operations for specific workflows. Each tool executes a pre-optimized query with explicit data access limits—the agent calls a resolved operation rather than reasoning over raw schema. This eliminates unintended data exposure, reduces token usage, and ensures deterministic execution across multi-step agentic workflows.

- Source Tools (coming Spring 2026) expose tightly defined operations specific to each system, mapping directly to approved actions and enforcing predictable execution, transactional safety, and auditability for production workflows.

Toolkits bundle a curated set of MCP tools into a single governed MCP endpoint, purpose-built for a specific team or use case. Organizations define exactly which connections, tools, and logic are exposed, and publish a single MCP Server URL that AI tools connect to directly. Each Toolkit can be deployed as a dedicated MCP server, ensuring agents operate only within their intended scope.

Query Federation Security

Connect AI enables queries that combine data from different sources directly, on-demand. When using federated queries:

- Permissions are checked for each source system in the federated query
- Users must have access to all data sources being joined
- Audit logs capture access to each source system involved

Audit Trails and Governance

Comprehensive audit trails are essential for compliance, security monitoring, and understanding how AI agents interact with your data. Connect AI provides full visibility and logging across every data query from AI.

What Gets Logged

Log Element	Description
User Identity	Authenticated user or agent identity making the request
Session Activity	Session start/end times and associated actions
Tool Invocations	Which MCP tools were called and with what parameters
Query Execution	Full query traces including source systems accessed
Timestamps	Precise timing for all operations
Error Reporting	Failed access attempts and permission denials

Compliance Monitoring

The audit trail capabilities support key compliance requirements:

- SOC 2: Evidence of access controls and security monitoring—demonstrate who accessed what data and when
- ISO 27001: Information security event logging—maintain records for security investigations
- GDPR: Data access transparency—document processing activities for personal data

Integration with Microsoft Agent 365

For organizations using Microsoft Copilot Studio, Connect AI integrates with Microsoft Agent 365 to extend governance controls:

- Centralized policy management across AI agents
- Unified audit trails combining Connect AI and Copilot activity
- Agent 365 tracing visibility into MCP tool invocations
- Consistent governance framework across all connected systems

Incident Response and Access Termination

When a security incident occurs—whether a compromised credential, anomalous agent behavior, or unauthorized access attempt—response speed and ability to respond matters. Connect AI provides multiple levels of access termination to enable rapid containment.

Immediate Termination Options

Connect AI offers granular kill switches at multiple levels, allowing security teams to contain incidents with precision:

Level	Action	Use Case
User	Revoke individual user access	Compromised account, terminated employee
Connection	Disable specific data source	Suspicious activity on one system
Workspace	Suspend workspace access	Isolate a team or use case
Account-wide	Emergency account lockdown	Major breach or critical incident

User Access Revocation: Individual user permissions can be revoked through the Connect AI admin console. When user access is removed:

- The user loses access to all assigned connections and workspaces
- CRUD operation permissions are immediately invalidated
- Subsequent query attempts return authorization errors

Connection-Level Disable: Administrators can disable individual data source connections.

Workspace Isolation: Workspaces can be suspended to immediately block all users assigned to that workspace from accessing bundled data sources.

SSO and Session Management

For organizations using SSO integration, access termination flows through your identity provider. Revoking user access in your IdP (Entra ID, Okta, etc.) prevents new authentication.

Recommendation: If your security policy requires immediate session termination, coordinate IdP revocation with direct user removal in Connect AI to ensure both new and existing sessions are blocked.

Audit Log Access During Incidents

During a security incident, rapid access to relevant logs is critical for investigation and remediation. Connect AI audit trails capture:

- User identity and authentication events
- Session start/end times
- Every tool invocation and query executed
- Data sources accessed and operations performed
- Failed access attempts and permission denials
- Precise timestamps for all actions

Incident Response Checklist

When responding to a security incident involving Connect AI:

- Identify scope: Determine if the incident affects a single user, connection, or broader access
- Contain immediately: Use the appropriate kill switch level to stop ongoing unauthorized access
- Preserve evidence: Export relevant audit logs before any changes that might affect log retention
- Investigate: Review audit trails for the affected timeframe to understand the incident scope
- Remediate: Address root cause (rotate credentials, update permissions, patch vulnerabilities)
- Restore: Re-enable access incrementally after confirming remediation

Security Review FAQ

Common questions from security reviewers and recommended responses:

Where is data stored?

Connect AI accesses data in place and never stores, copies, or caches customer data. All queries execute in real-time against source systems. Data sovereignty requirements are inherently satisfied because data never leaves its original location.

How is authentication handled?

Connect AI supports SSO integration with major identity providers and protocols (SAML 2.0, OpenID Connect, Microsoft Entra ID, Okta, etc.) and implements OAuth 2.1 with PKCE for secure authorization flows. User Credentials mode ensures individual accountability when accessing data sources.

What compliance certifications do you have?

CData has completed third-party SOC 2 Type II and ISO/IEC 27001:2022 compliance audits. The full security documentation package including audit reports is available upon request.

Do you conduct penetration testing?

Yes. CData engages independent third-party penetration testers for regular security assessments of Connect AI, in addition to ongoing internal security testing between external audits. Third-party assessment summaries are available under NDA.

How do you prevent unauthorized access?

Connect AI uses RBAC passthrough to honor existing source system permissions. Additional layers include CRUD operation scoping, connection-level access controls, and derived views to pre-filter sensitive data. Users only see what they're authorized to access in the source system.

What encryption is used?

AES-256 encryption at rest – AES-128 is the minimum level for Federal government compliance – and TLS 1.3 in transit. All credentials are stored encrypted and never exposed in logs or responses.

How are audit logs maintained?

Every data interaction is logged with user identity, session activity, tool invocations, query execution details, timestamps, and error reporting. Logs support SOC 2, ISO 27001, and GDPR compliance requirements.

What about prompt injection risks?

Connect AI implements multiple safeguards: RBAC passthrough ensures even if an agent is manipulated, it cannot access data the user isn't authorized to see. CRUD scoping limits destructive operations. Derived views pre-filter sensitive data before it reaches the agent context.

Why use a managed platform vs. self-hosted MCP servers?

Self-hosted MCP servers introduce supply chain risks (unvetted code), lack standardized security controls, require DIY authentication and audit logging, and create distributed management overhead. Connect AI provides centralized security, governance, and compliance as a managed service.

Implementation Checklist

Use this checklist when deploying Connect AI to ensure your implementation follows security best practices.

Pre-Deployment

- Document data sources that will be connected and their sensitivity levels
- Identify user groups and their required access levels
- Review existing RBAC configurations in source systems
- Determine SSO provider and authentication requirements
- Establish audit log retention and review policies

Authentication Configuration

- Enable SSO integration with your identity provider
- Configure User Credentials mode for supported data sources
- Verify OAuth/PKCE flows are working correctly
- Test RBAC passthrough with representative user accounts
- Configure SCIM 2.0 provisioning with your identity provider to automate user lifecycle management

Access Control Setup

- Configure user permissions per data source
- Apply CRUD operation scoping (default to read-only unless write is required)
- Create workspaces for different use cases/teams
- Define derived views to pre-filter sensitive data

Governance and Monitoring

- Configure audit log settings and retention
- Establish process for regular audit log review
- Set up alerts for anomalous access patterns
- Document incident response procedures for security events

Ongoing Operations

- Schedule regular access reviews and permission audits
- Monitor for new user onboarding and offboarding
- Review and update derived views as data requirements change
- Stay current with Connect AI security updates and best practices

Summary

CData Connect AI provides the enterprise security infrastructure required for data connectivity to support production AI deployments. By leveraging the platform's built-in security features—SSO integration, RBAC passthrough, SCIM provisioning, CRUD scoping, Customs tools and toolkits, workspaces, derived views, and comprehensive audit trails—IT leaders can confidently deploy AI agents while maintaining control over data access and compliance requirements.

Key Takeaways

- Managed platform eliminates supply chain risks of unvetted MCP servers
- In-place access means data never moves—source system controls remain authoritative
- Identity-first security through RBAC passthrough—no need to rebuild permission structures
- Layered access controls—combine source permissions with Connect AI's CRUD scoping
- Data curation with workspaces and derived views—expose only what's needed
- Comprehensive audit trails—full visibility for compliance and security monitoring
- On-Premises data connectivity under the same governed layer with the Connect Gateway

Request Your Security Kit

CData provides a comprehensive security kit for enterprise evaluations, including SOC 2 Type II reports, ISO/IEC 27001:2022 certification documentation, and detailed architecture diagrams. Contact your CData account team or visit cdata.com/security to request the full security documentation package.