Ask the Expert: Architecture Q&A Session

Pooja Kachawaha, Community Events & User Groups Coordinator, Boomi

boomi

Upcoming Community Events

Ask the Expert: Architecture

- **Date:** February 3, 2026
- **Time:** 1:00 PM EST

Join us for Boomi Office Hours – your go-to Q&A session led by Boomi Professional Services experts!

Higher Education User Group Meetup

- **Date:** February 11, 2026
- **Time:** 1:00 PM EST

Join us as we meet for our monthly virtual Higher Education User Group!

Ask the Expert: Integration Development

- **Date:** March 3, 2026
- **Time:** 1:00 PM EST

Join us for Boomi Office Hours, a live Q&A session with Boomi Professional Services experts!

https://usergroups.boomi.com/

Ben Shepherd

Senior Systems Engineer
Boomi

- Platform Architect in Boomi Advisory Services
- Technical Architect in the PSO APJ Team
- 20+ year integration experience
- Mentor for Boomi PSO APJ Peers

Questions

Question:

Architecture & Templates

- What standardized architecture templates are available for greenfield Boomi projects?
- Is there a centralized repository where these templates and best practices are maintained?

Load & Performance Testing

- What load and performance testing strategies are recommended for Boomi integrations?
- How can these tests be effectively executed in pre-production environments to closely reflect real-world system behavior?

Architecture & Templates

Greenfield Boomi Projects

First step - Know the products

Second step - Know the business objective. Why was Boomi a fit and where do we want the product/s positioned?

Third step - Know the integration backlog

Fourth step - Know what the expertise level of your business

Load & Performance Testing

Greenfield Boomi Projects

Load Testing

Establish a baseline

Get familiar with the expected load in the current environment. This will allow you to determine how to set up threads for HTTP Requests.

Prepare your integrations/processes

Following Boomi development best practices, you would become familiar with the facade concept.

Know how you want your tool to call Boomi processes

Boomi Atomsphere API Execute or HTTP Request to API Gateway or direct to Shared Web Server.

Use a tool that provides visuals/analytics to quick peaks, latencies and error conditions, such as JMeter.

Load & Performance Testing

Greenfield Boomi Projects

Performance Testing
Once you have identified potential bottlenecks, inspect the Boomi process to get familiar to potential bottleneck positions using Process State within Boomi Process Reporting.
You can also utilise Boomi supplied dashboards that allow you to inspect the performance of API calls.

Question:

Is Boomi MCP gonna be tackled?

MCP (Currently in Tech Preview)

Short Answer = Yes (MCP Server and MCP Client)

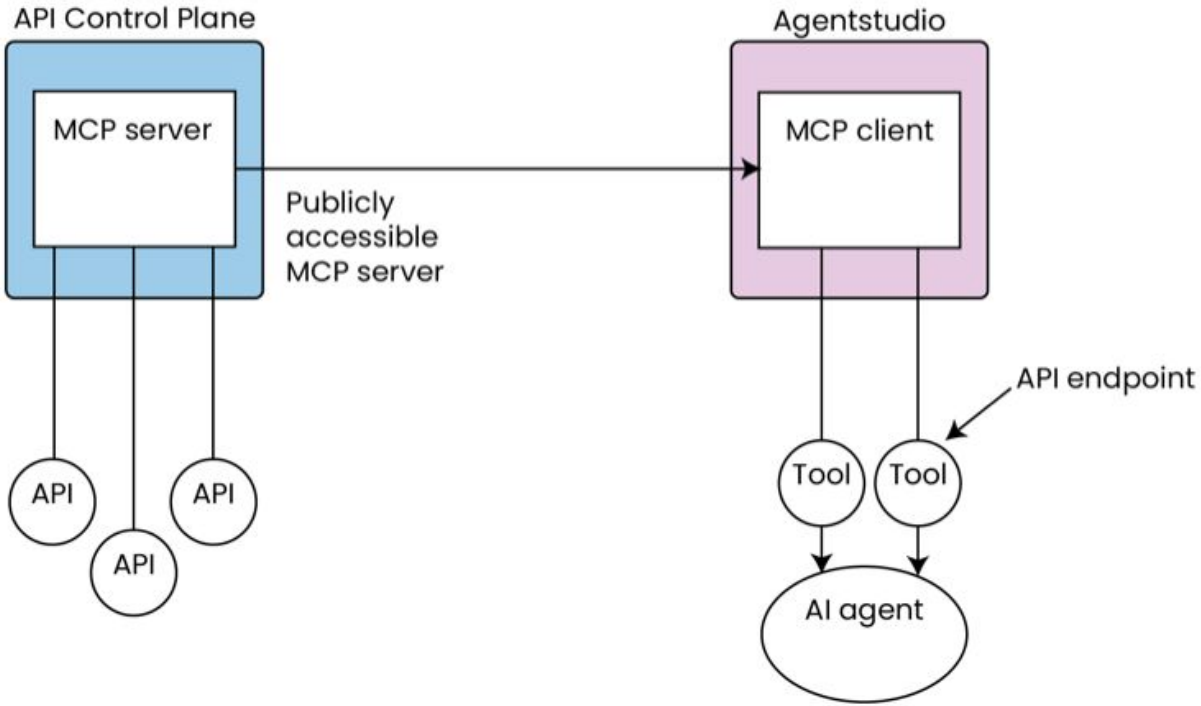

Question:

After deploying Atom Watch across production and non-production Boomi Molecules, we are leveraging Boomi HTTP logs as our primary signal for API-level observability (traffic, latency, error rates), while process logs remain available for internal execution detail. In your experience, what is the recommended observability model for Boomi-based APIs: HTTP-layer-first with selective process log correlation, or a more blended approach? What trade-offs should teams be aware of at scale?

APIM Observability Model

Open Telemetry

List of Vendors: Vendors | OpenTelemetry

Traces

Metrics

Logs

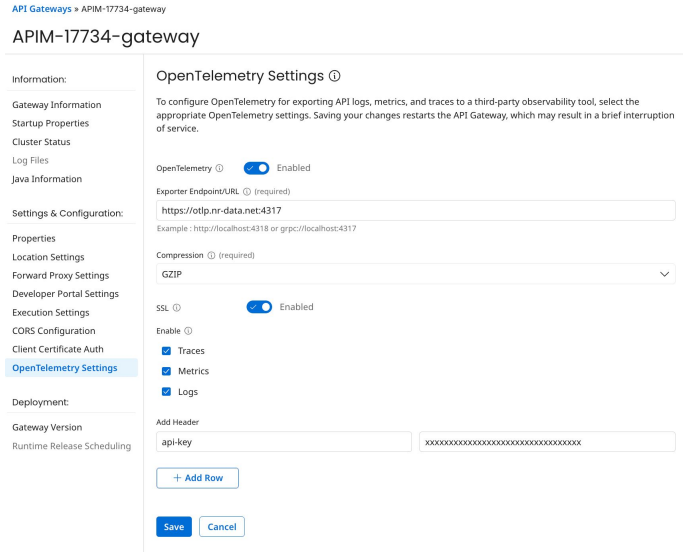

APIM Observability Model

Open Telemetry with DataDog

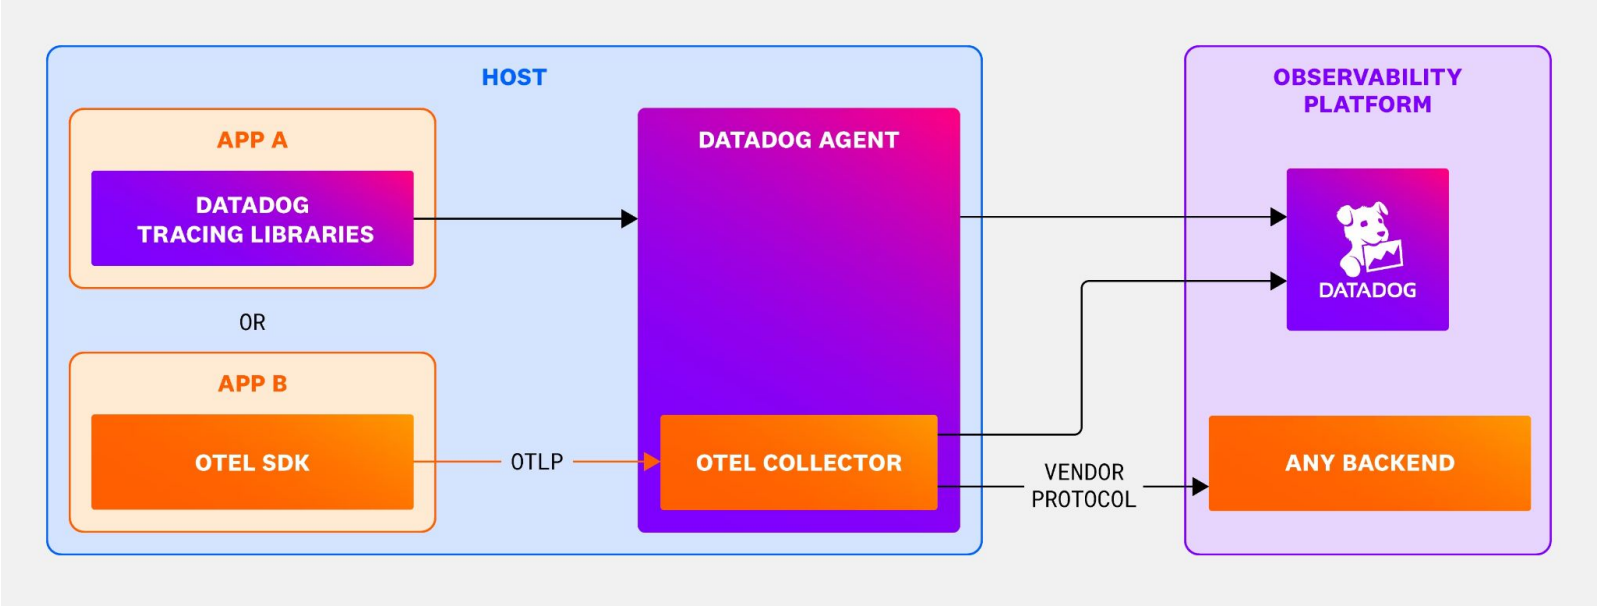

Question:

Best practices for error handling and test automation

Boomi Best Practices

Error Handling

Error Handling is typically considered to be :

- guaranteed delivery
- Service Level Agreement (SLA) for certain error scenarios.
- Aggregation (Avoiding Error Storms)
- Tracking Identification (multiple executions)
- Reusable

Customers can choose Advanced Frameworks as part of their Platform Blueprint to have the Error Handling Notification Framework setup and configured.

Boomi Best Practices

Test Automation

First of all, you need to come up with a set of reusable error scenarios with predicted inputs/outputs that can be used to perform assertions.

What assertions are valid for the test scenario.

What to automate:

- Unit Tests
- Integration Tests
- Connection Tests

You could create a Test Suite, that calls all of these tests and package up and deploy to your API Gateway to allow for use within the CI/CD pipeline.

Boomi Best Practices

Test Automation (cont)

Option: You can use the Boomi Marketplace to use the Automated IUTG (Integration Unit Test Generator) Agent:

The Automated IUTG Agent creates unit test case scenarios by examining process metadata. It takes in a process ID and generates sample input and expected output for that process. This agent is a headless agent that can be used in tandem with an Integration process to test developed processes.

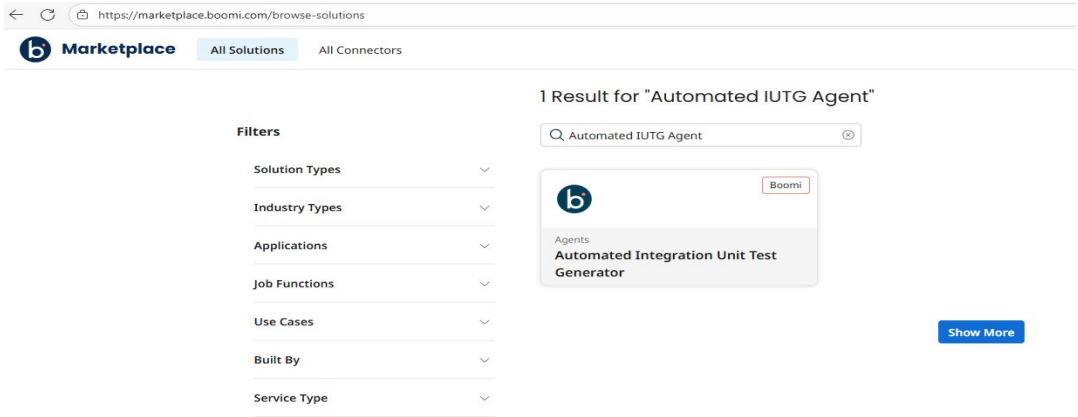

Question:

How does one project certain agents from being used by certain users? i.e. I have an HR user who uses an agent for onboarding, and I have an employee user using an agent for booking leave, but the employee user should not be allowed to use the HR agent for onboarding

AI Agent Access

HR versus Employee user access

At the moment there is no Agent specific user access but it is in the roadmap.

If this is the situation then it is better to have another layer on top of the agent, either Flow based application or the API managing the Agents and using Agent connection in the integration processes.

Question:

With Agentic the new Buzz word . How do we re-architecture traditional integration patterns and design to meet or ever know these APIS or integrations would be part of some Agentic AI use case if ever .

Integration Patterns in Agentic AI

Role of APIs

- RAGs optimise output LLM and references an authoritative knowledge base before generating a response (Eg: ChatBot). It is reactive.

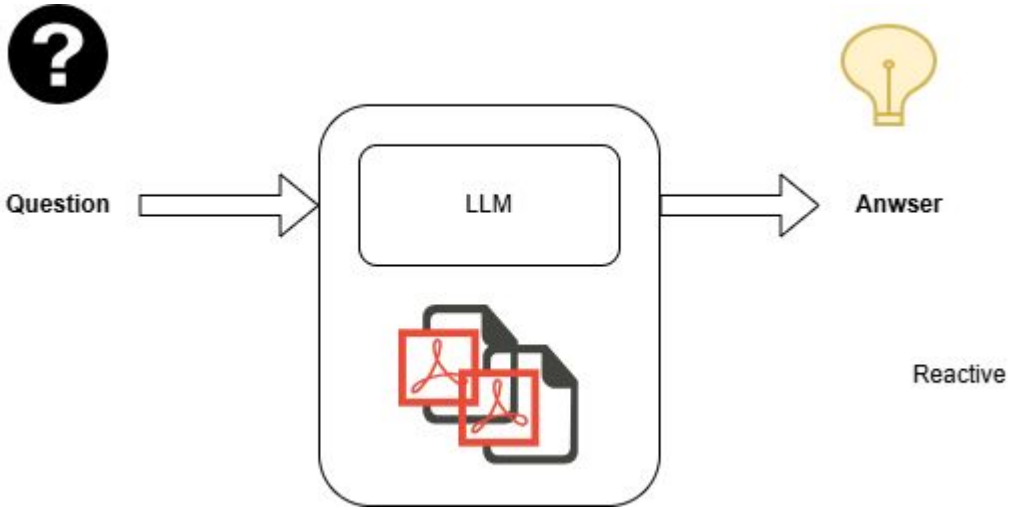

Integration Patterns in Agentic AI

Role of APIs

Let’s have a look at an example what makes up an Agentic AI Solution:

- Goal Oriented Planning
- Multistep Reasoning
- Autonomous Decision Making
- Access to Tools, Knowledge and Memory

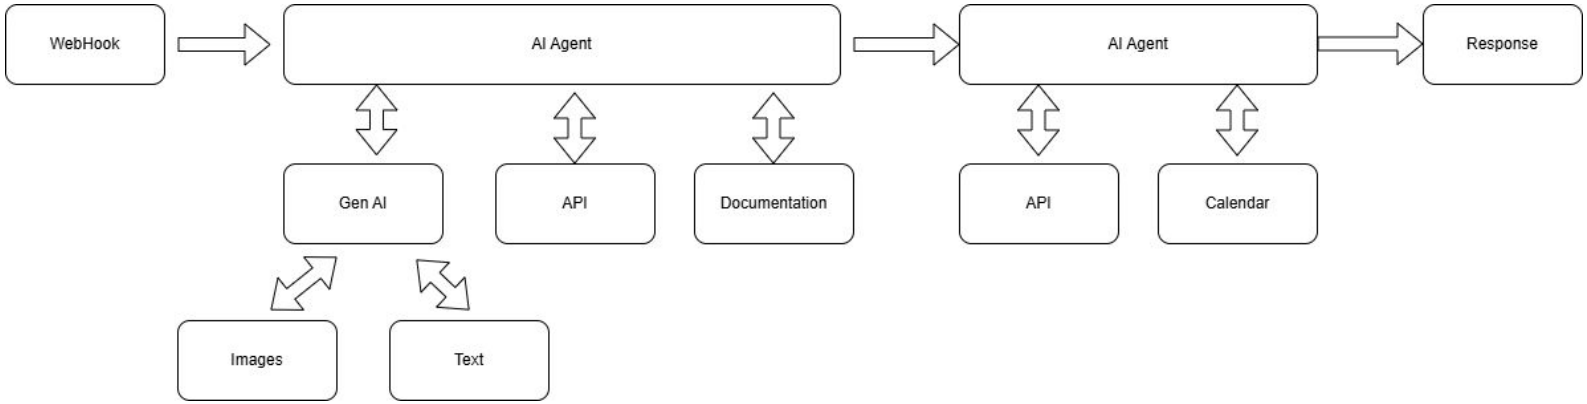

Integration Patterns in Agentic AI

Role of APIs

So in terms of API Integration Patterns, the philosophy is the same, however, in terms of establishing the plan to achieve your goals, you may need to consider the behaviour of your AI Agents.

The design of an Agentic AI is limitless.

Question:

What are the few major points of caution while designing an integration ecosystem during this current AI wave?

Cautions with Designing Integration during AI Wave

Major concerns and Protection of assets

In majority of causes the integration ecosystem can remain the same.

The main concern would be the the designers and developers of AI Agents should consider the use of certain API services.

Main concern - Data

Protection

- Limit what is returned from API hosted.
- Limit the data sent to 3rd party AI models.

All Boomi AI Agents should consider the Establishment of guardrails.

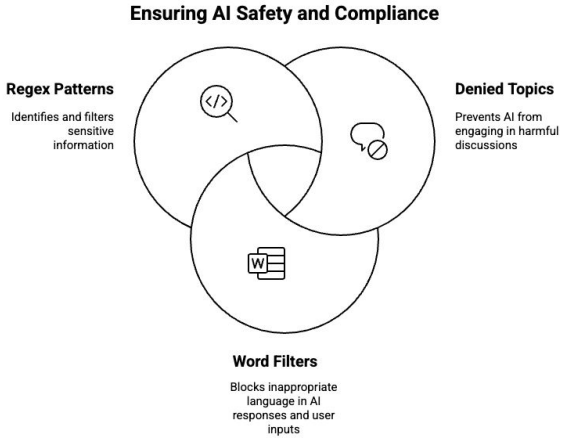

Question:

What is the best Practice in using Event Streams instead of Parallel processing?

Event Streams Opposed to Parallel Processing

Best Practices

Parallel Processing

- Ideal for running threads to process multiple executions simultaneously managed by the flow control shape.

Event Streams

- Ability to queue messages in a topic outside of the runtime to execute processes based on subscribers to the topic.

Would use Parallel processing to optimises the processing of bulk data by processing execution across multiple threads.
Would use Event Streams to decouple the processed in the case of Guaranteed Delivery or Pub/Sub requirements.

Question:

I wanted to know the best practices for implementing the integration processes

Integration Development

Best Practices

Leverage a governance document

Establish naming conventions

Establish integration patterns

Ensure detailed design is approved before development

Consider reuse of components

Use extensions

Question:

New updates

New Updates

January Release

Boomi has many new updates this month. Best to look at:

https://help.boomi.com/docs/Atomsphere/Release_Notes

To see what is related to your account.

New Updates

API Management

We added these features

- Added the following new roles and privilege:
  - **API Administrator** role: This new role grants full read and write access to all API Management pages, including the Dashboard. The role eliminates the need for users to create custom roles to achieve full administrative access. It consolidates all read-write privileges into a single role, replicating the full capabilities of the **API Management - Access** privilege while ensuring permissions are clearly defined. (APIM-14256)
  - **API Support** role: This role provides read-only access to all API Management pages. The role enables support personnel to view system resources and access the APIM Dashboard. It explicitly includes the new **API Management - Gateway Logs Download** privilege, allowing users to download gateway logs for debugging purposes. (APIM-14255)
  - **API Management- Gateway Logs Download** privilege: This privilege grants users the ability to download logs directly from the Gateway Information page, facilitating streamlined debugging and issue resolution. This enhancement enables support personnel to retrieve critical diagnostic information without requiring "Write" access to the Gateway configuration, provided they hold the **API Management - Gateway Read** privilege to access the page. (APIM-18472)

New Updates

API Management

This release upgrades the H2 database and provides a seamless auto data migration path for existing gateway. (APIM-18627)

- **We fixed this issue**
- When querying a gateway in GraphQL, the `classification` for environments was not being returned. (ACP-2448)

New Updates

Integration

We added these features in December

- Added execution timeout configuration options to the Agent step configuration component so that processes no longer wait indefinitely for a response from the agent. (INT-16451)

We added these features in January

- The **OK** button is now visible in the Environment Extensions dialog only to users who have edit access to the environment extensions. (INT-15948)
- Updated the validation message for creating global variables, providing clearer guidance that names cannot start with a number. (INT-16500)
- Renamed the **Variable Assignment** tab to **Value Assignment** for global variables (INT-16286).
- Added a referencing dropdown for non-encrypted text fields, allowing you to reference runtime configuration values by typing the global variable syntax. (INT-16092)

New Updates

Integration

We added these features in January (cont)

- The ExecutionRecord API returns a 503 error if the service is down instead of returning zero documents with a success code. If you try to get the execution record in Process Reporting, a 503 error is returned as well. (INT-16259)
- Added polling for OAuth2 AccessToken generation API, so you can track the status of the OAuth request flow and receive the access token upon completion. (INT-16111)
- You can specify a timeout value for Agent Step configuration components, ensuring processes fail gracefully rather than hanging indefinitely when external services are unresponsive. (INT-16451)
- A new **Stop If No Documents** option is available under the **Process Options** tab to prevent execution of downstream steps when no documents are generated, improving process efficiency. This setting applies only to the individual process and is not inherited by sub-processes, which you must configure separately to enable the same behavior. (INT-15091)
- The recommendation pop-up in the **Start** shape side panel prompts you to enable **Stop If No Documents** when the option is unchecked, improving process handling. (INT-16041)
- Process execution has been enhanced to halt at steps that generate no documents, avoiding downstream execution and minimizing exceptions during test mode and process runs. (INT-257)
- Added user-friendly logging for steps that produce no documents, enabling you to understand why a process execution was halted or skipped at a particular step. (INT-15945).

Question:

How to balance process decoupling for a customer whose license model is consumption based?

Consumption Model

Balance Process Decoupling

The consumption model is volume based

Process Decoupling is ideal for async processes

Balancing customers volume, could be done by setting up limits.

API Gateway have counters in a plan. The plan can reject requests that go over a quota. Integration have counters / or process properties as well which could be used to provide limits.

Question:

I want to know while learning about Boomi what are the best approaches we can do like what procedure /structure /roadmap /way we follow to get the best fit.

Starting at Boomi

Best practices

Boomi provide a wealth of knowledge articles that go through best practices, but agree, they is a lot of information to absorb.

Ideally, you could reach out to Boomi and/or Boomi partners that can support you in your initial setup and you can go from there.

There are also online training as well as Boomi Instructor Lead training that can get you up to speed at a faster rate.

Question:

What recommendations does Boomi have for sizing a runtime atom or instances of a runtime molecule considering CPU, memory, and minimum network bandwidth. What metrics would be used to validate either down scaling or upscaling instances?

Runtime Sizing

Basic and Clustered Runtimes

During Boomi PSO Platform establishment, the Technical Architecture tends to ask questions to get familiar with the metrics, budget, application landscape and other business requirements.

Both of these pieces of information support the decision on what runtimes, number of environments and runtime type that best suits the customers needs. The sizing can be done my reviewing the metrics and determining the size (CPU, RAM) required but also the amount of space the file system needs to handle the data loads.

Ofcourse, each decision would expect some performance testing to go to ensure that the setup is a good fit.

Question:

I'm curious about your best practices for troubleshooting and debugging, particularly 1) seeing the current value of a variable without adding notify or message shapes, 2) seeing the actual request body sent to an API endpoint, 3) testing a listener-driven integration process.

Support Troubleshooting

Best practices

Looking at variables in flight within a process without Notify and Message shape is a challenge as logging typically describes the shape being executed and the duration of the shape being executed in the container logs.

If you want properties being audited, consider running in test mode and looking at the log during a set properties shape that appear in the platform logs.

Seeing the body being sent to the API request is something that is non-standard due to amount of data being rewritten to disk. Again, running in test mode is an idea where you can check data in request in the REST connector shape or other API connector shape.

For listeners, this is tough as typically we would not suggest logging for these styles of processes.
For all three of these situational troubleshooting error points may have data being feed to an audit/error logging. This would allow you to feed data to logs outside of the main workflow.

Question:

What is boomi doing to close the gap in API Management Systems that offer more modern features, such as GraphQL, Websockets, MCP etc?

Boomi API Management

Types of service - GraphQL

GraphQL

Boomi accept GraphQL by accepting tradition GraphQL queries which present themselves in JSON format.

https://help.boomi.com/docs/Atomsphere/API%20Management/Topics/t-api-Leveraging_your_GraphQL_API_Subscription

Boomi allow expose GraphQL to call the Platform
Primary Endpoint:

```
https://api.boomi.com/graphql
```

Boomi API Management

Types of service - WebSockets

Web Sockets

Boomi has managed web sockets be putting in front a proxy to convert wss to https.

This being said, you may have a TCP Connector built for Execution and Listening.

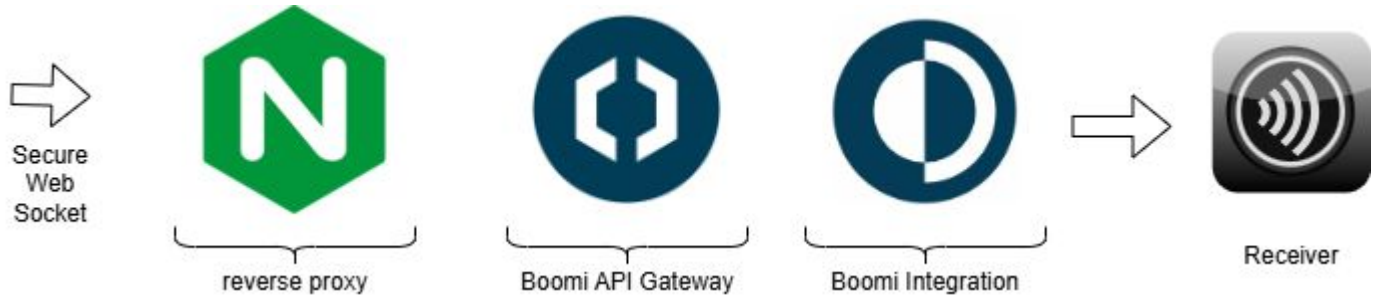

Boomi API Management

Types of service - MCP Server and Client

MCP

https://help.boomi.com/docs/Atomsphere/Platform/MCP_overview

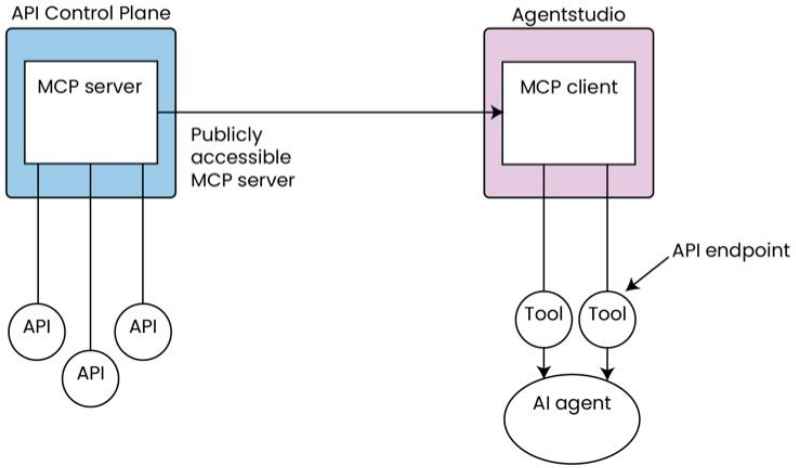

MCP (Currently in Tech Preview)

Sequence of interactions between MCP Client and Server

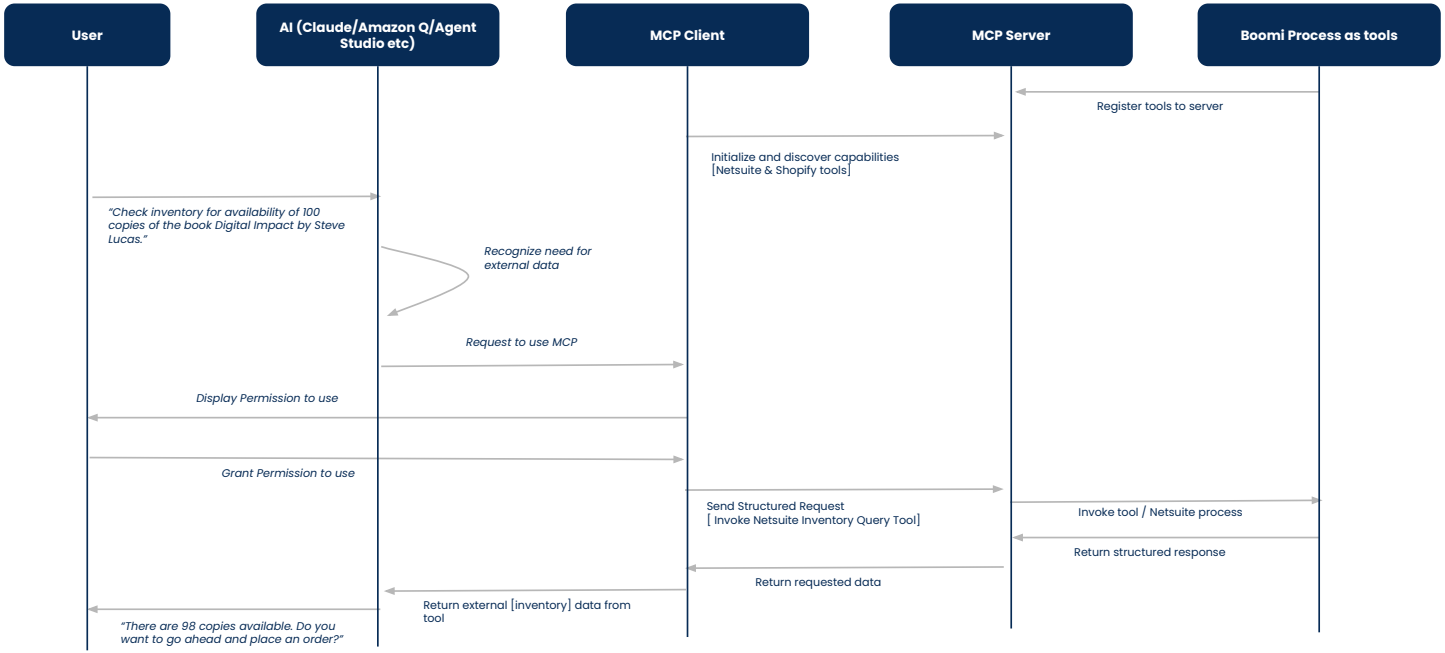

Question:

How boomi will support for large data volume, complex workflow and custom scripting not supported by pre-built connectors

Large Data Volume / Complex Workflows / Custom Scripting

How will Boomi support?

Boomi have products to support customers with Large Data Volumes by providing the ability to allow both horizontal and vertical scaling on multiple platforms, whether on on-premise or in the cloud.
Let’s consider large files, Boomi has MFT and Data Integration for ETL processing. The products allow option to perform parallel process or external processing using Event Streams, complex workflows can be decoupled into logic workflows and custom scripting allow for imported libraries where third party developers could support in the usage of their libraries. The product also allow the use of internal and external cache options to allow for memory management.
The Boomi has many partners that provide support, Boomi support allows for customer s to raise tickets and the Boomi user community is a resource rich space where users contribute every day.
Boomi also offers professional services and Advisory services to support customers.

Question:

I have a Boomi Gateway and a Runtime Cluster with 3 nodes. Do I need a load balancer between the Gateway and Runtime Cluster to distribute incoming API requests to the appropriate cluster node? Or is the distribution of works are to be handled by the Head Node in the cluster?

API to Integration Landscape

Load Balancing

Ideally, you need 2 load balancers (NLB and ALB)

Although, Boomi Integration has an inbuilt LB, let’s not confuse with the need for an external LB for ingress traffic.

Questions

Thank you