# Ask the Expert: Architecture Q&A Session

Pooja Kachawaha, Community Events & User Groups Coordinator, Boomi

boomi

# Upcoming Community Events

## Meet Data Connector Agent: Build Any Custom Connector in Minutes

- **Date:** July 8, 2026
- **Time:** 10AM EDT | 3PM BST

Join us for Boomi Data Integration: DCA Community Webinar!

## Higher Education User Group Virtual Meetup

- **Date:** July 15, 2026
- **Time:** 11AM PDT | 2PM EDT

We invite you to our monthly Higher Education User Group meeting, a dedicated platform for professionals in the education sector to connect, share, and learn.

## Activating MCP With Boomi

- **Date:** July 22, 2026
- **Time:** 10:30AM IST | 3PM AEST

Join us to discover how Boomi's enterprise-ready MCP implementation can connect your AI agents to the systems that power your business.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

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

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Questions

**Question:**

Integration best practices

# Integration Best Practices

## CoE Governance

Some Best Practices are:

- Leverage a governance document (Development Guidelines)
- Establish naming conventions
- Establish integration patterns
- Ensure detailed design is approved before development
- Consider reuse of components
- Use extensions
- Error handling Framework

Refer to Article: Boomi Blueprint: Comprehensive Guide - Boomi Community

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

1. Molecule best practices around the Private Atom Cloud Architecture.
2. Best practices around setting the correct Heap Size, Memory Utilisation tips etc

# Boomi Cloud Runtimes
# FKA Private Atom Cloud

There are now two types of Boomi Cloud Runtimes that have replaced the legacy Private Atom Cloud.

**DCS (Dedicated Cloud Service)** which is an elastic cloud runtime.

Single-tenant
Boomi managed
Self-service configured

**MCS (Managed Cloud Service)** which is managed by Boomi and has some flexibility, but non elastic.

Single-tenant
Fully managed
Custom sized
Advisory included

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud Runtimes Best Practices

## Number of Execution Workers

Knowing that each execution worker is a JVM, you need to consider how many execution workers would be required to manage the predicted loads.

Knowing that each execution worker have a heap of 512 MB, the RAM allocated to the runtime plays a vital role when selecting the number of execution workers. Also note the maximum number of concurrent executions is 50.

So you would also need to consider the number of system threads to handle the processing of data at peak/burst periods.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud Runtimes Best Practices

### Correct Heap Size

MCS manages this for you, typically we start with 512MB and consider 1024MB for larger clouds.

Assume that DCS follows the same approach.
Key consideration: based on concurrent executions and expected document size.

If there is a need for larger heap size due to OOM conditions, then you can raise a ticket for Boomi to make the appropriate setting.

Need to consider Non-Heap memory (system level memory pressure).

Optimal performance requires **balanced JVM heap size**, coordinated thread management, and process-level memory optimization while monitoring key runtime metrics.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud Runtimes Best Practices

## Memory Utilisation

### We first look at what uses memory.

Within the main JVM heap, the biggest consumers are:
- document size multiplied by concurrency (a 10 MB XML document held across 5 process steps by 20 concurrent executions commits roughly 1 GB of old-gen heap)
- map function caches and cross-reference lookups held for the lifetime of each execution
- compiled process component definitions cached at startup (hundreds of deployed processes create significant permanent overhead)
- Groovy script engine data if resetScriptEngineData is set to false
- and connector caches for replay IDs, CSRF tokens, and lookup caches from active listener processes.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud Runtimes Best Practices

Memory Utilisation (cont)

Looking at large documents - Split documents early and process one at a time for batch processing.

Set Xms equal to Xmx in atom.vmoptions. If Xms is lower, the JVM expands the heap on demand at runtime, and each expansion triggers a Full GC.

Lower com.boomi.container.transform.maxCacheSize. The default is high and controls how many document elements are held in heap during mapping. Reducing it for complex large-document transformations significantly lowers peak old-gen usage.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

- Why limited for number of molecules - Current version of java libs in boomi core vs impacts (old java vs vulnerabilities) - Migration to the newer java version 17/21 (when) - Java versions vs build new connectors (limited usage in Java 11) - Third party libraries - Old groovy version 2.4 vs last groovy version - Old vs new APIM (see Boomi APIM vs Boomi Cloud APIM) - etc.

# Boomi Recommendations

## Clustered Runtime

The Clustered Runtime (FKA Molecule) has a recommended range of between 3 to 10 nodes.

Reason:
- Minimum of 3 is due to HA requirements (during rolling restart)
- Maximum of 10 due to the stabilization of the runtime.
  - Head node comm to secondary nodes (Hazelcast-based cluster discovery (multicast/unicast on port 45588/7800) and JGroups coordination
  - Shared filesystem convention (File locking)
  - Forked JVM expansion (50 forked JVMs per node can lead to 500 at 512MB each = 256 GB). Any mode node would be unmanageable.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Recommendations

## Java Version 8 and 11

The limitation of Java 11 is known to Boomi Product team, but they also consider other factors like connector requirements that may need

In terms of security support, refer to Amazon Corretto | endoflife.date v17 and v21 now have an EOL expectancy less than v8.

**Note:** when Boomi elect next LTS version, it would need to support Java 8 libraries.

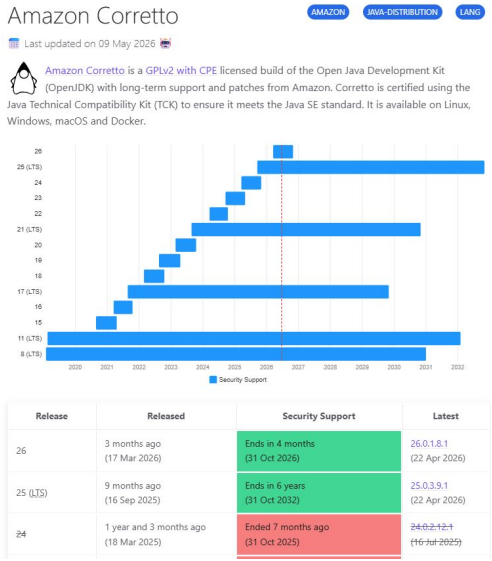


Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Recommendations

## Groovy 2.4

Groovy 2.4 is the recommended version of Groovy for the Boomi runtime.

The latest version of Groovy is version 4, which has a much richer set of libraries compared to the version 2.4 that is currently used in Boomi. While you could potentially upgrade to version 4, it may disrupt their existing scripts and processes with Boomi as there are some deprecated classes in v2.4.

Groovy 4 is currently not on the roadmap, however, it is definitely a future consideration.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Recommendations

## Old and New APIM

Old Boomi APIM

- Gateway on Atom/Molecule
- Single-tenant, on-premise
- Basic OAuth 2.0, API key
- Limited policy engine
- No multi-gateway federation
- No shadow API discovery
- Manual API documentation
- Tightly coupled to Integration runtime

New Boomi APIM (2025+)

- Cloud-native elastic gateway
- Multi-cloud + on-premise
- Full OAuth 2.0, OIDC, mTLS
- Dynamic policy enforcement
- API Control Plane: any gateway
- Shadow / zombie API discovery
- AI-generated docs (Boomi Scribe)
- MCP server exposure + governance
- Rate limits, quotas, traffic mgmt

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

What are best practices when it comes to utilising parameters in some of the API connectors? I have had some issues where I set parameters, but it overrides the input document, which seems to override the point of setting a parameter (e.g., an ID at the end of an endpoint). Notably for POST connectors

# Use of Parameters in HTTP POST

## Correct Use

Using the REST Connector, the pattern for a dynamic ID in the path is:

- In the operation's Options tab, set the Path to your base resource path, e.g. `/api/v1/orders/`
- On the connector shape on the canvas, open the **Dynamic Operation Properties** tab (not the Parameters tab)
- Set the `Path` property, sourced from a DDP like `DDP_SYSTEM_ORDER_ID`, which resolves to `/api/v1/orders/12345`

Because this overrides the URL path and nothing else, your POST body flows through entirely untouched.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

How do companies usually deal with error handling? We would like to get aggregated information on errors and warnings, as well as be able to send errors from integration to business users when appropriate, ideally in a unified way so that a stream of errors can be handled across all integrations.

# Error Handling Framework

## Concepts

Error Handling for a customer depends on the support requirements for the organisation. Many customers have their own strategies, however, there are commonalities.

1 - Error Repository (Disk, Database)
2 - Decoupling from workflow
3 - Error profiles
4 - SLA
5 - Email alerting

For accumulation or errors in email, you can do this within a process (puopulting a document cache) or from a central data repository that can be scanned periodically to collect all errors and report.

For a centralised error repository, you could build a dashboard.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

How to avoid/minimise Garbage Collector full pass time? This lead to VIEW_MISMATCH and a rolling started is initiaed at cluster level.

# Garbage Collection

VIEW_MISMATCH Events

The root cause chain is important to understand clearly before touching anything:

A Full GC stop-the-world pause freezes the JVM completely, which means Hazelcast heartbeats stop. Once the pause exceeds the heartbeat timeout threshold, the other nodes eject the frozen node from the cluster view. When the frozen node wakes up and tries to rejoin, the membership views no longer match — that's the VIEW_MISMATCH. Boomi's rolling restart is the correct recovery, not the problem itself.

Start with considering the GC algorithm that works for you.
G1GC (default <16GB heap), ZGC (Java 11+, large heaps), Shenandoah (Java 11+, if G1GC not returning heap to OS) and CMS (Avoid)

Heap Sizing Rules:
Xms = Xmx
Leave 2-4 GB for the OS and NFS I/O Buffer
Bigger heap = Longer Full GC

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

When can we see spec driven engineering via Boomi AI to aide in migration

# Boomi Migration using AI

## Spec Driven Development using AI

Boomi Companion is an AI tool that accepts prompts to build integrations based on criteria. Idea would be to feed in a spec to build an integration, but there are many factors to building an integration.

For migration, you would need to consider a few things:
1 - What is the current solution business logic.
2 - Is the integration idempotent
3 - What is the messaging pattern. Async or Sync.
4 - What are the messaging profiles and mapping requirements
5 - Will the integration need to feed into Data Hub
6 - Who needs to know when an error occurs and how will the integration align with the error handling framework

So, for these concepts, you can create references within the AI project to support your migration development

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

Boomi architecture best practices and design principles

# Best Practices

## Boomi Architecture and Development

- Native First, customer scripting last
- Modular over Monolithic
- Document Tracking
- Consider Process Mode for each Process
- Separation of components across process layers (Orchestration, Business Logic, System, Framework)
- Idempotency
- Clean Code (no orphaned shapes or unterminated branches)
- Use environment extensions (appropriately) on all parent processes

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Best Practices

## Folders

- Keep Connectors in a central location (preferably at the top of the Component Explorer)
- Maintain a folder structure for your processes aligned with the Development Guidelines set by your Integration Architect
- Avoid nesting folders past 4 levels deep.
- Never have Sandbox connectors in the main integration folder system.
- Consider a Shared folder for components to be shared across the team not just individuals.
- Use folder level permission (Boomi roles) to restrict write access from non developers.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Best Practices

## Naming Conventions and Process Designs

- In your Governance Document (or Development Guidelines) consider setting up a naming convention for all the developers to align with.
- Consider Process Design Standards
  - Process Mode
  - Reusable connectors
  - Integration Patterns
  - Environment Extension consistency
  - Error Handling

Use Boomi Blueprint for a list of Boomi defined Development Guidelines: Article: Boomi Blueprint: Development Guidelines - Boomi Community

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

We are trying to implement solution with JH core with other system, need to know about any kind data security in case of client wants to use Boomi Cloud.

# Data Security

## Boomi Cloud Runtimes

### To confirm JH Core is Banking/Investment System

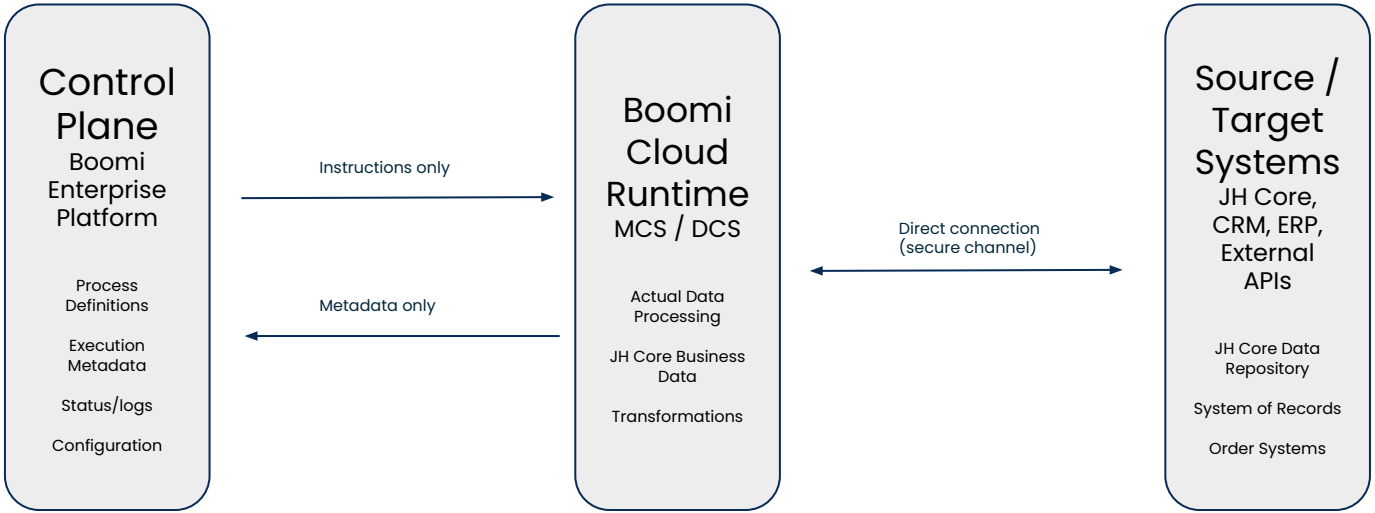


Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Data Security

## Boomi Cloud Runtimes

### Dedicated Cloud Service (DCS)

Integrations with less-sensitive data where VPN to on-premise is still needed

Single-tenant — not shared with other customers. VPN connectivity to on-premise JH core banking. Boomi manages the infrastructure but data flows through Boomi's dedicated environment.

### Managed Cloud Service (MCS)

Enterprises wanting full Boomi management with private connectivity and compliance advisory

Single-tenant, VPN/VPC peering, highest management support. Suitable if bank has contractual data processing agreements with Boomi covering banking data

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Data Security

## Boomi Cloud Runtimes

### Compliance

Boomi is CPS 230 compliant. This compliance applies to Banks, Insurers and superannuation trustees and places specific obligations on institutions to identify, assess, and manage risks arising from their service providers—including technology providers like Boomi—especially when those providers are considered material to critical operations.

Reference: https://boomi.com/wp-content/uploads/CPS-230-Mapping-Guide.pdf

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

How can we enable AI to build a good architecture framework which can handle generic error and after fixing those error Boomi process can be rerun without user or developer intervention.

# Using AI to Build Best Practices

## Boomi Companion

My suggestion here is to use Boomi companion. The AI tool already builds Boomi processes. However, if you want to build good architecture, then you need to define what this is to your tool.

Like most AI Agents, the Agent is supported by the toolset assigned. You can build plugins that support your tool to allow more definitive reasoning to prompts.

Example: I want to enforce that only a maximum of 20 shapes on a canvas (where possible without impeding the business functionality), then you set that as a guardrails for the AI Agent. Each prompt will then consider the rule whenever making a decision on how to build the process.

I suggest start small and then introduce more rules to align with the customer architecture principles.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

Is it better to go with Listener Process or platform event for real time Synchronization. Advantages & disadvantages of both listener and platform event

# Real Time Events

## Listener versus Platform Event

Let’s explore the two options:

Listener - Designed for real time events. If you have a web services listener, then you can allow for synchronous processes. For queue based, this is asynchronous, however, you have guaranteed delivery within the Boomi space. Be sure to consider Boomi Event Streams or other for this as this allows the outsourcing of memory and storage from the Boomi runtime. Also, have DLQ to capture failed documents.

Listeners are considered to typically have sub-second response time.

Platform Event - Ideal for asynchronous integration where you want to manage the documents/event outside of Boomi. An important note is to ensure you have a replay Id assigned to the event so that Boomi can capture and replay to allow for guaranteed delivery.

Platform Events are considered to typically have 2 to 3 second response time.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

I am new to this platform and soon want to use this in integration and migration project so what should be the necessary things I should know and how can I make out the best use out of this platform as an developer.

# Getting Starting

## Boomi Platform

Ideally, you should first get familiar with what the Platform can do. I recommend starting Boomi training (I did this when first starting train.boomi.com)

1 - Platform Fundamentals
2 - Developer 1, 2 and 3 (Associate Developer)
3 - Professional Developer (Developer)

This will give you understanding on how to use the shapes in the processes to best achieve your goals.

Explore other courses regularly to keep to to date.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Getting Starting

## Boomi Platform

When you have the knowledge, consider the approach to your projects. What does you current landscape look like and what are the functions?

Consider the right tools for the job:

1 - APIs consider Boomi APIM and Integration Runtime
2 - What are the TPS for the integrations? This would allow you to plan how you Boomi platform is architected to achieve business requirements.
3 - Migration may require data sync (CDC) etc, you may want to consider Data Integration.
4 - Large File transfers or SFTP to SFTP integrations. You may want to consider Boomi MFT.

Idea here is to focus on performance and longevity of your architecture and not put additional load on Integration Runtime if can be avoided.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

**Question:**

Boomi CAM: Call transformation, custom adapter

# Boomi Cloud API Management (CAM)

## Call Transformation and Custom Adapter

Call Transformation allows CAM to plug in pre-processors.

Idea is that an endpoint use can call using a certain protocol, say SOAP when REST is required, the CAM Traffic Controller can translate the request  from SOAP to the required REST protocol.

The Transformation can allow be used for post-processors.

Reference:

Call Transformations | Boomi Documentation

Configuring Call Transformation for an Endpoint | Boomi Documentation

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud API Management (CAM)

## Call Transformation and Custom Adapter

For custom adapters, the developer will require Local Edition SDK from Boomi. The SDK requires JDK 11.

The Adapter class  implements the CAM SDK processor interface. Key methods are:
1 - preProcess(), and
2 - postProcess()

Both methods receive pre-input and post-input JSON configuration defined in the Control Center Call Transformation Page.
Note: Any failures MUST result is a call failure (not silent passthrough)

Compile and package the jar (**note:** nested directories not supported in Traffic Manager) and name customer-extension.zip, copy to assemblies directory and run build-images.sh
This results in the adapter being visible in the Call Transformation configuration as a Processing Adapter option.

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Boomi Cloud API Management (CAM)

## Call Transformation and Custom Adapter (Use Case)

Create CAMs API Definition and connect it's traffic manager to an external Backend API. ACP can discover the API Products and can call the processes.

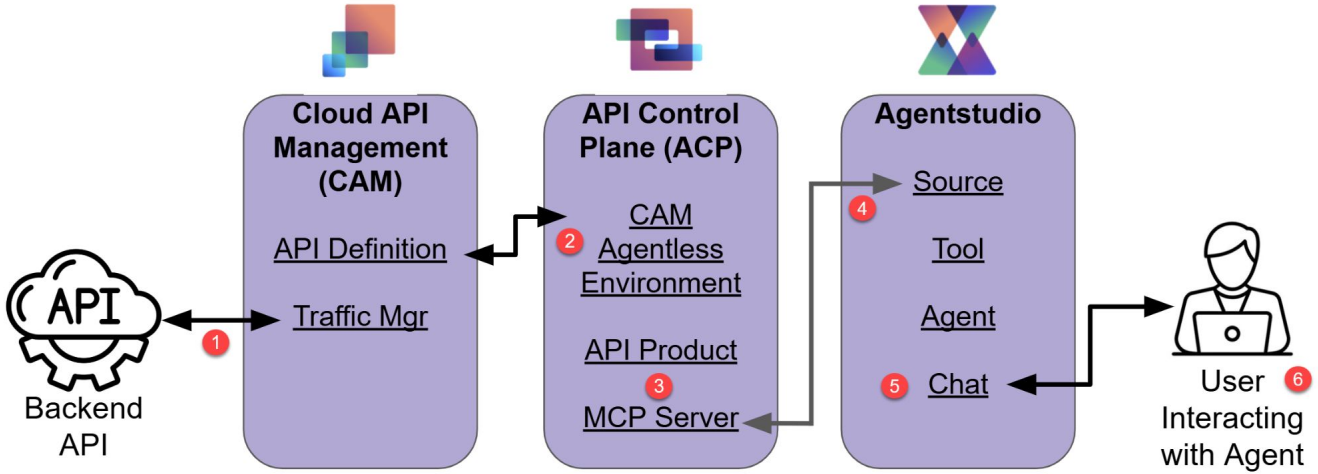


Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.

# Questions

# Thank you

Copyright © 2025 Boomi, LP or its affiliates.
All rights reserved.