Ask the Expert: Architecture Q&A Session

Regan Knepp, Community Events & User Groups Lead, Boomi

boomi

Upcoming October Events

Higher Education User Group Meetup

Date: October 9, 2024, 2:00pm EDT
Registration: Please Register here
Agenda: Join us as we meet for our monthly virtual Higher Education User Group!

Transportation & Logistics User Group Meetup

Date: October 15, 2024, 1:00pm EDT
Registration: Please Register here
Agenda: Join us for a virtual transportation & logistics focused user group event, where industry professionals come together to share insights, gain expert knowledge, and build valuable connections.

Document and Process Properties

Date: October 23, 2024, 11:30am EDT
Registration: Please Register here
Agenda: What is a Boomi Document? What's the difference between Document Properties and Process Properties, and how do I use them? Join Senior Integration Captain as he explains these details, talks through use cases, and reveals some tips and tricks.

https://usergroups.boomi.com/

Jim Dines

Technical Architect

- 5 years at Boomi as a Technical Architect.
- Expertise in architecting all Boomi runtime platforms (Atom, Molecule, Private Atom Cloud) on-premise and on cloud platforms like Azure and AWS.
- Skilled in deploying Boomi on Kubernetes.
- Proficient in resolving performance issues, integration architecture and design, and using the Atomsphere API.
- Experienced in implementing CI/CD with Boomi and managing API integrations using Boomi API Gateways.
- 35 years of experience in IT across Manufacturing, Life Sciences, and IT Consulting industries.

Questions

Question:

How to create a Java library and embed into Boomi then use Groovy to call the functions inside this library.

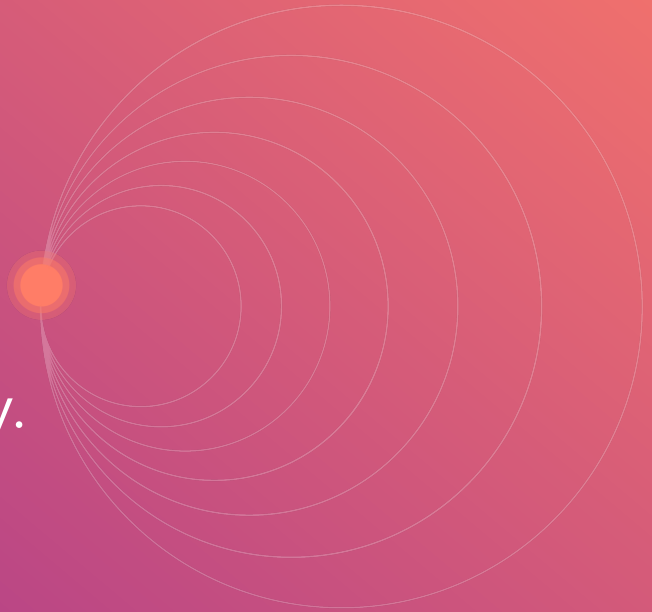

Boomi Custom Scripting

Utilizing JARs

There is a pretty good user guide article for this.

https://help.boomi.com/docs/Atomsphere/Integration/Process%20building/c-atm-Custom_scripting_35340a9d-a236-403a-9f53-1563fceb61a5

The steps are:

- Create Your JAR file
- Upload JAR files into your account library (Settings > Account Information and Setup > Account Libraries)
- Restart the runtime
- In your groovy script, import the objects in the JAR. For example… import com.my.custom.jar.MyClass;

Question:

Is there anything different we need to do for deploying Boomi on Kubernetes?

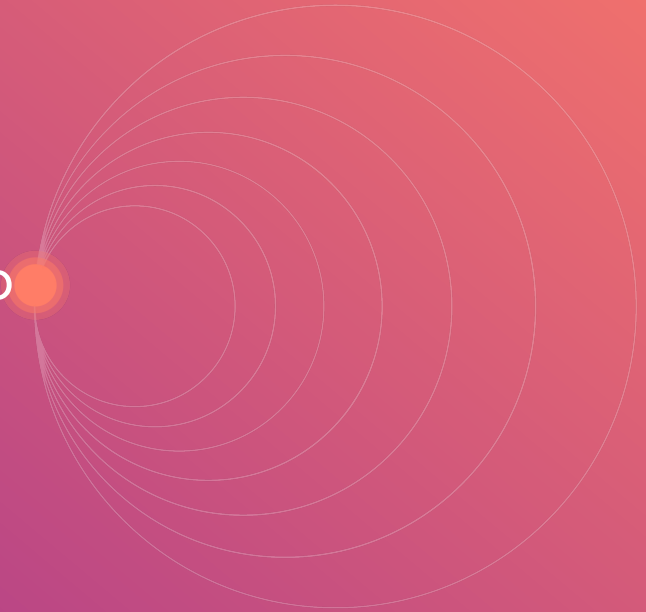

Boomi on Kubernetes

Differences on K8s

From a licensing POV, there is no difference..

There are many other differences, here is just a few…:

- K8s is a platform that virtualizes applications. You will need running K8s clusters.
- Your cluster will need to support Ingress and Services. This usually implies a load balancer local to the K8s cluster and an Ingress controller, like the Nginx ingress controller.
- You will want to determine how you will network your K8s cluster. You will typically use either Kubenet or a CNI (Container Network Interface).
- You will want to carefully plan how you will secure your cluster's and your application's networking using both Ingress and Services.
- Boomi runtimes run in containers in Kubernetes. For Clouds and Molecules, and NFS is still required to allow the Boomi runtime to persist.
- You can implement autoscaling of Molecules and Clouds in K8s. Even without autoscaling, scaling runtimes is very easy in K8s.
- Deploying Boomi runtimes in K8s is done by executing either YAML or HELM scripts that drive the installation. Boomi can provide YAML scripts.
- You can use DevOps platforms to execute the creation of runtimes by configuring pipelines that run the installation scripts. You will need to decide if you want to trek down this path.

Question:

What are some strategies for Runtime optimization?

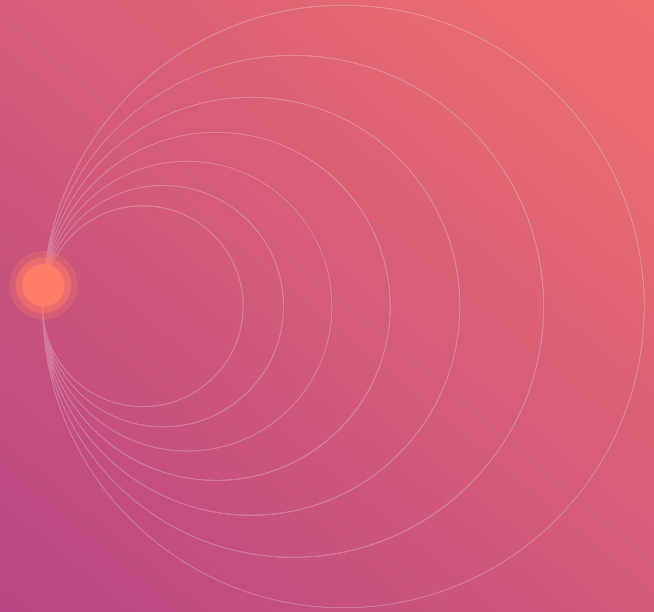

Runtime Optimization

Strategies

The first thing you want to do to manage performance and capacity is to monitor your Boomi runtimes

- Monitor both JMX parameters (heap sizes, queue lengths, concurrent executions, etc…) and system parameters (memory, cpu, NFS IOPS, etc….)
- Create Alerts - Alert on key parameters at warning and critical thresholds
- Develop a discipline around reviewing key parameters and making adjustments as needed based on trends.

  ** Boomi PSO can provide documentation on monitoring parameters, alerting thresholds and in most cases, templates to use with various APM systems.

Some basics:

- At least initially, heaps should be sized to be ½ of available memory. Heap sizes should not exceed 12GB
- For low throughput runtimes 2 cores is adequate, for medium to high 4 cores should be used, for high volume runtimes, 6 or more may be required.
- Use the property, Maximum Simultaneous Executions per Node, to control how many executions can run simultaneously per node. This can help minimize the chance of overloading heap and CPU. You can then monitor the number of queued processes and the average execution time to detect changes in demand over time.

Question:

How to use Queue in Boomi?

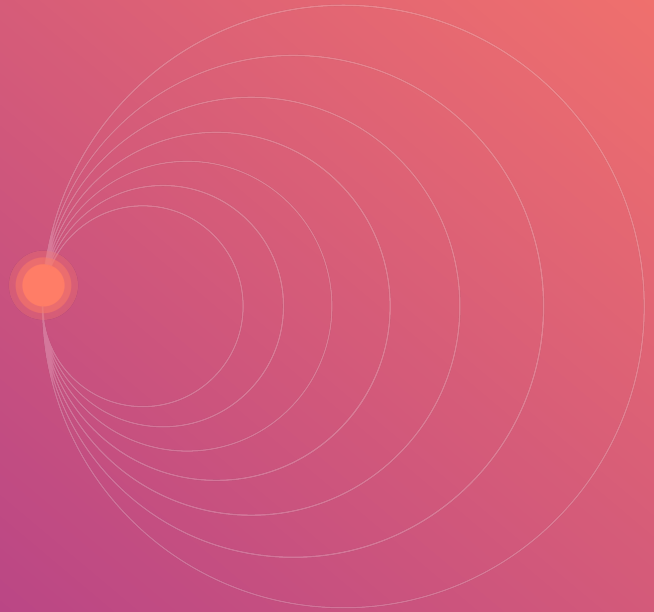

Atom Queue Patterns

Inherent Patterns

- Point to Point - Queue is used between a single producer and a single consumer
- Pub Sub - Queue is used between a single producer and multiple ocnsumers

Queues are useful for enabling..

- Resilience - Queues can be used to buffer data so that data loss is minimized due to process failure
- Reuse - Queues can be used to enable many producing processes to utilized the same running process
- Scalability - Queues can be used to create reuse patterns that increase scalability by utilizing systems resources more efficiently

Other Techniques

Decoupling, Process Reuse, Fire and Forget

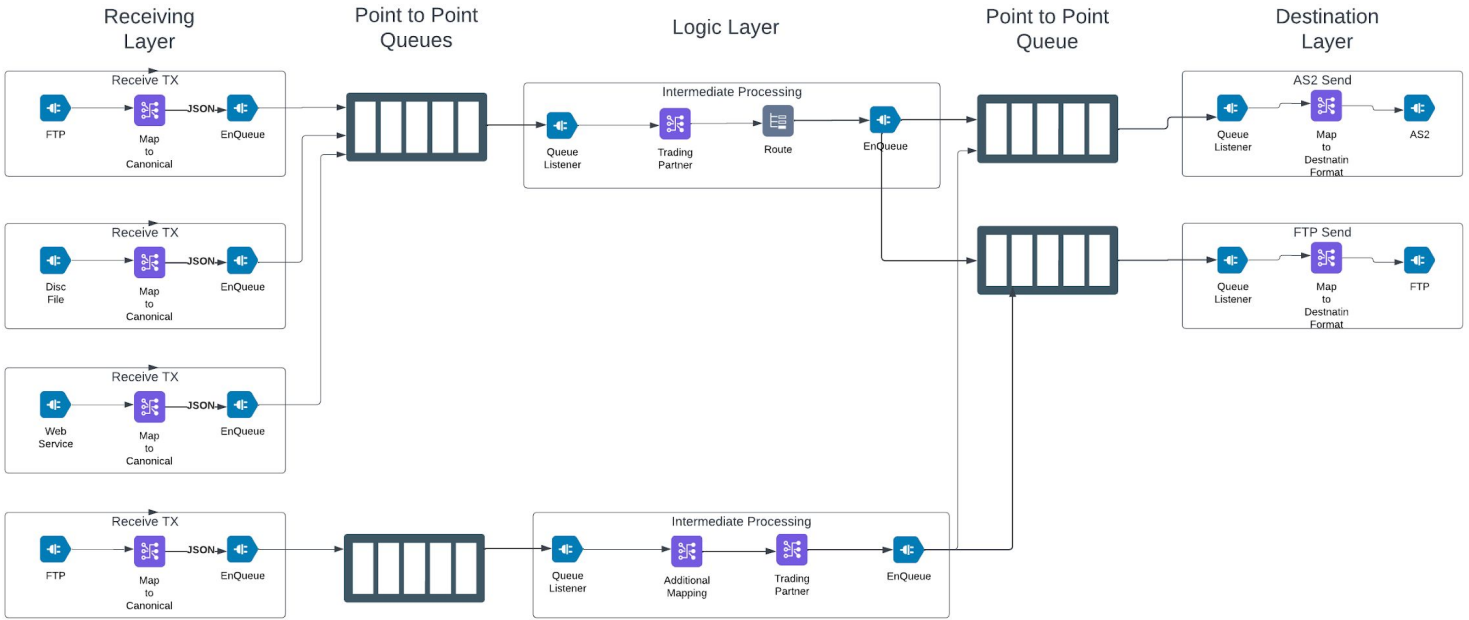

Other Techniques

Decoupling, Process Reuse, Fire and Forget

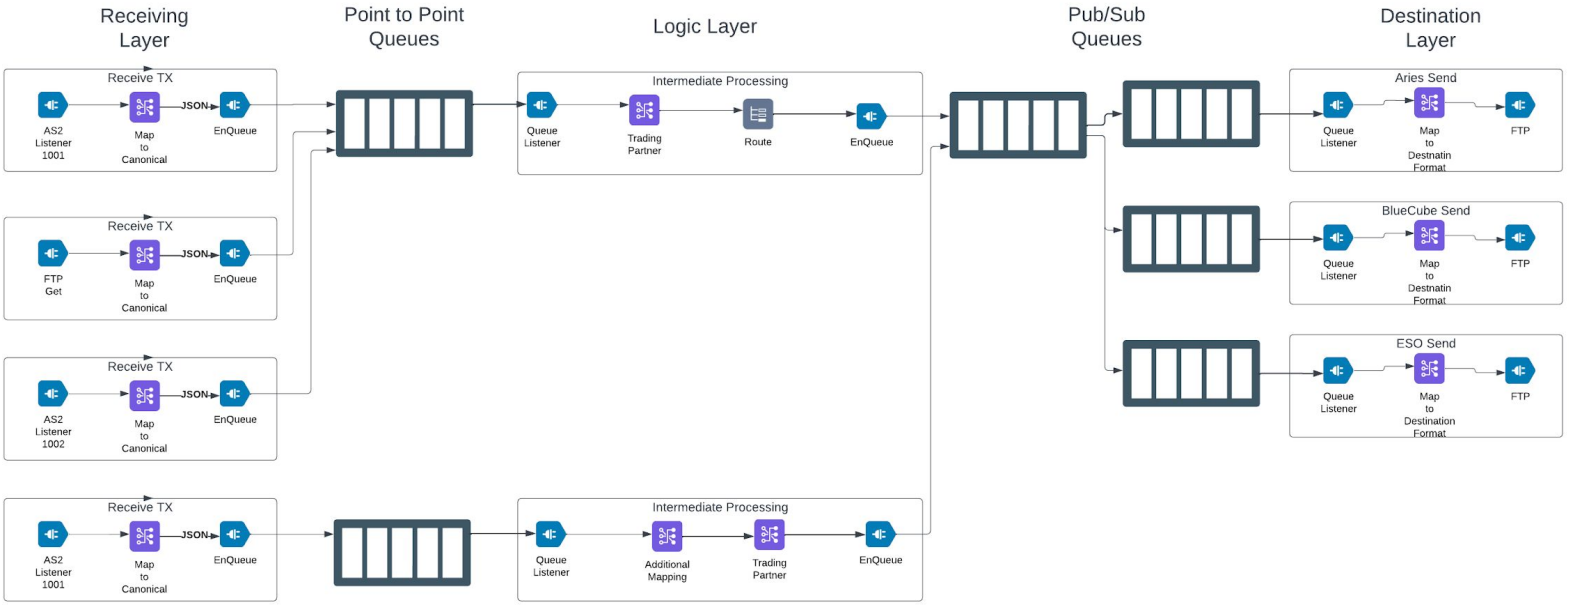

References on Atom Queues

User Guide Article on Atom Messaging

https://help.boomi.com/docs/atomsphere/integration/event-based%20integration/c-atm-atom_message_queuing_5310fd40-efdf-4bb2-bb3c-6ef099472b36/

Boomi Booster on Atoms Message Queues

https://community.boomi.com/s/article/Working-with-Boomi-Atom-Message-Queue

Question:

Security when hosting in Azure.

How to allow restart of atoms without access to atom management?

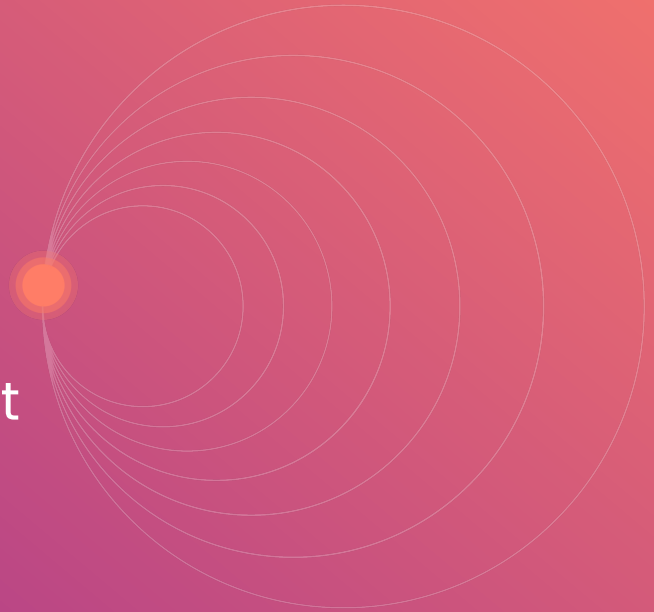

Boomi Security in Azure

There are several things to consider when deploying Boomi into Azure.

SSO - For Boomi platform logins, you can integrate into Entra ID. This is fairly straightforward integration. You will need to set up an Application for Boomi in Entra ID, upload a certificate and you will need to add users or groups to the associated Enterprise Application. You can then populate Boomi with the required Authorization Endpoint, Federation ID name and certificate.

Similarly, for **Oauth2 used with APIs**, you can set up applications and use Entra as your Identity provider.

Runtimes exposed to the internet should be protected with NSGs and a WAF. A common architecture for exposed runtimes would be to use an Application Gateway for load balancing the runtime and Use Front Door with a WAF enabled in front of the App Gateway.

How to allow restart of atoms without access to atom management?

This is tricky……

To do this you will have to build a component external to the Boomi Platform.

The following article explains how to do this from a Boomi process…

https://community.boomi.com/s/article/schedule-a-molecule-rolling-restart-using-custom-scripting

If you are executing the runtime as a service as recommended by Boomi, you can execute the commands to start and stop the services on each node in the runtime. For example, on Linux, use systemctl restart <service_name>.

If you are not running Boomi as a service, you can use the atom restart command on each node.

Question:

Recommendations for shared storage architecture - server-based NFS is generally a single point of failure, but Boomi has ruled out Spectrum Scale (GPFS), Gluster, NFS Ganesha, etc.

Any experience with NetApp or other NAS solutions?

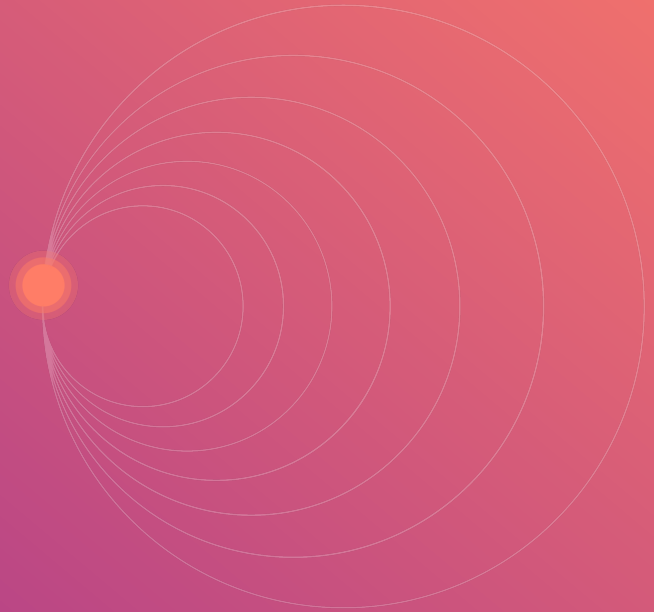

Shared Storage for Boomi Clusters

Platforms

The vast majority of customers we see are using the following successfully..

On Premise - Most customers are using a storage subsystem that provides NFS or SMB volumes. Common ones are... Dell Isilon, NetApp and Pure Storage... For heavily loaded systems Flash storage in these platforms is of huge benefit.

Azure - Azure Netapp Files, Ultra capacity pools.

AWS - EFS

We also see a few active/passive NFS servers built on infrastructure using Pacemaker and DRDB. This usually requires some configuration in the NFS server to increase the number of NFS daemon processes. The default is usually 8-16. The number of daemons needs to be much higher than that. More like 128.

Shared Storage for Boomi Clusters

IOPS and Throughput

The real key is to understand the available IOPS on the NFS volume. IOPS is the key performance measurement for the NFS. Using the fio command to test the IOPS will help understand available throughput.

$ cd /mnt/nfsvol

$ mkdir fiotest

$ cd fiotest

$ sudo apt-get install -y fio

$ fio --name=BoomiTest4k --rw=randwrite --direct=1 --ioengine=libaio --bs=4k --numjobs=16 --iodepth=128 --size=100M --runtime=600 --group_reporting

With that fio test, you should see a minimum of around 14,000 IOPS. As demand grows in your Boomi runtime, you may need as much as 100-120,000 IOPS for a very heavily utilized, large cluster.

Question:

What are the main advantages of Cloud Agent versus Local Agent?

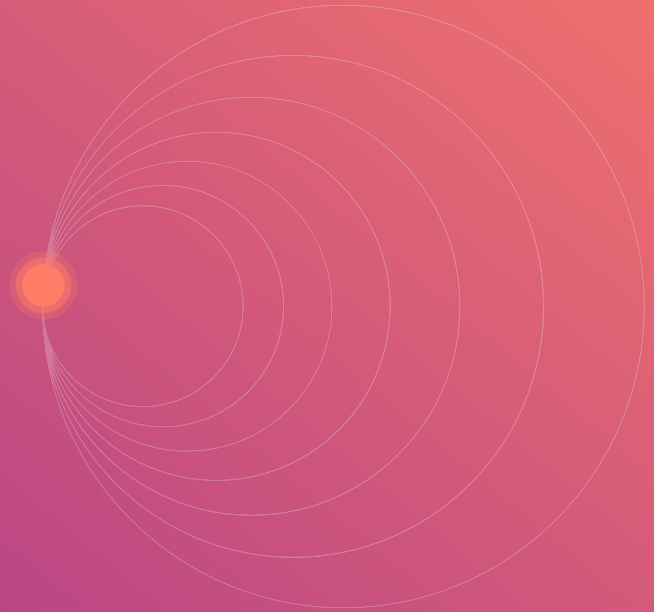

Private Cloud vs On Prem

Other than satisfaction of corporate standards or financial goals, the primary advantage of a cloud implementation vs a local implementation are:

- Local implementations have direct, secure access to locally deployed endpoint systems as well as SaaS and other external endpoint systems.
- Cloud deployments are easier to build and managing infrastructure. Access to endpoint systems on local, on premises networks usually requires VPNs or other mechanisms.

Question:

What are the most useful industry tools for monitoring the health of Molecules and Gateways?

What are your experiences using Azure DevOps in a CICD pipeline? We are in the infancy of this research and looking for pros, cons, and challenges.

Tools for Monitoring

Just about and APM (Application Performance Management) platform will do a good job of exposing all of the performance aspects of a Boomi runtime. The ones we commonly see are:

- DataDog
- New Relic
- Azure Application Insights
- Dynatrace
- …. There are many more

Key features to look for are:

- The ability to monitor JMX. Boomi related parameters are exposed via JMX
- Dashboarding
- Alerting

Boomi PSO Azure DevOps Experience

We frequently work with customers to help them implement pipelines in Azure DevOps

We have produced an integration layer that enables DevOps tools to manage deployments in Boomi.

We don’t provide Azure DevOps services, per say. Unless it is directly related implementing CICD workflows for Boomi.

We have consulted at the Architecture level on CICD in DevOps and Boomi Runtime Deployments on AKS using Azure DevOps.

Questions

THANK YOU!!

boomi