

Digibee Platform Technical

Threshold and Limits

Revision History	2
Overview	3
Growing Up your Technical Limits	3
Quotas	3
Technical Limits Definitions	4
Processing Capacity	4
Egress Traffic	5
Message Rates and Queueing	5
Message	5
EPS (Executions Per Second)	5
Queueing	6
Execution Log Retention	6
Completed Execution Logs	6
Pipeline Logs	6
Digibee Storage	7
Object Store	7
Relationship Management	8
Connectivity	8
VPNs and ZTNA	8
Technical Usage Limits	9
Test Environment Usage Limits	9
Production Environment Usage Limits	10
Usage Limits for All Environments Combined	10

Revision History

Version	Change's description (Section, Page or text)	Date
1.0	Document creation and publication	April 4, 2022
2.0	Document Update	February 15, 2024

Overview

Technical Usage Limit means the limits of the Digibee Platform resources allocated for your realm. These limits, also known as quotas, are applied to make sure Digibee Platform is fairly used by all SaaS realms to protect the overall health of Digibee's SaaS ecosystem.

Depending on the number of Pipeline Subscriptions, Test or Production RTU Capacity, your Realm will account for a specific usage limit. These limits include, but are not limited to, how many users, pipeline deploys, infrastructure resources and dedicated infrastructure you have on the Digibee Platform.

This document details what are the technical usage limits for the Digibee Platform.

Growing Up your Technical Limits

There are two ways to scale up your limits as described below:

- **Proportional:** proportional limits are the ones that apply in direct proportion to the number of RTUs in each test or production environment. They grow as a linear function of the number of RTUs.
- **Range:** range limits are the ones that apply to specific ranges of RTUs in each test or production environment. They do not grow linearly but in blocks.

Quotas

Quotas are applied in two different scopes:

- **Contracted Quotas:** it includes all pipelines that belong to your entire realm. Your total quotas depend on the size of your contract with Digibee.
- **Deployed Quotas:** it is applied to each deployed pipeline. In this case, your total quota depends on the size of the deployment. For example, the memory or CPU for a pipeline deployment (included in a RTU) cannot be shared with another deployment.

Technical Limits Definitions

Processing Capacity

RTU contains 64MB of memory and 20% of a modern CPU core for the pipeline to execute. It also provides 10 processing threads for concurrent flow execution.

In order to scale processing capacity, the platform allows users to do it in two different and complementary ways:

- **vertical scale:** means adding more processing capacity to a single instance of flow execution
- **horizontal scale:** means adding more copies of an instance of flow execution

Vertical scale (ie. pipeline deployment size)

The platform allows pipelines to be deployed in multiple pipeline sizes which will consume more RTUs. Currently there are three pipeline deployment sizes available in the Digibee platform: small, medium and large.

The following table shows how the processing capacity grows in each of the deployment sizes:

Deployment size	CPU ¹	Memory	Concurrent Executions	RTUs
Small	20%	64MB	10	1
Medium	40%	128MB	20	2
Large	80%	256MB	40	4

¹ typically a modern CPU core provided by processors such as Intel Xeon.

The defined memory and CPU allocations must accommodate the pipeline flow structure (ie, all the components, their relationships, capsules, etc), underlying services that help the pipeline execute its tasks and the payloads that are to be processed by the pipeline. Defining the exact pipeline deployment size for a given integration is a task to be performed during the pipeline design and initial tests. Every pipeline flow has its own particularities that might require more capacity.

Horizontal scale (ie, replicas)

Each pipeline can be deployed in multiple replicas, which basically adds more horizontal capacity to the pipeline. Each replica is a complete copy of a pipeline deployment. That said, if a pipeline is deployed as small as 1 replica, that would account for 1 RTU. Similarly if the pipeline is deployed as small size with 2 replicas that would be 2 RTUs. The overall formula for RTU calculation is:

RTUs = # RTUs for the pipeline deployment size X # replicas

Egress Traffic

Egress traffic accounts for all the traffic that leaves the Digibee SaaS Platform. This might include, but is not limited to: calls to external APIs, systems, protocols that are not hosted inside the Digibee Platform, VPN connections, use of any component or capsule that makes calls to external providers.

Message Rates and Queueing

Message

A message is the payload processed whenever a pipeline receives a request to process it, performs the processing and produces a response.

EPS (Executions Per Second)

EPS is the metric which calculates how many messages are processed by a given pipeline in a single 1-second window.

The maximum number of EPS is calculated for the whole Realm. As an example, if a Realm contains 20 RTUs for a production environment and the amount of EPS per RTU is 1, then the total amount of EPS supported by the Realm is 20 EPS. This number is shared across all deployed pipelines.

Note that this is different from the maximum concurrent executions allowed in a specific pipeline deployment size. As an example, a pipeline with a maximum of 10 concurrent

executions will process 1 execution per second when each message is processed with an average processing time of 10 seconds. Conversely, the same pipeline would process 10 executions per second when each message is processed with an average processing time of 1 second. So EPS is directly correlated with the average processing time and not necessarily the maximum number of concurrent executions.

Queueing

The platform allows pipeline queues to fill up with pending messages to be processed.

The quota imposed on queueing is calculated per all RTUs deployed on a single pipeline. As an example, if a pipeline is deployed with 4 RTUs then that pipeline queue will support up to 10 Gb of messages in the queue at any given time. Calculation: 4 RTUs x 2.5 Gb per RTU = 10 Gb.

Execution Log Retention

The platform tracks the execution of every single message processed by the system. Messages are retained according to time, number of messages and amount of bytes. It is displayed in the Monitor area of the platform under Completed Execution or Pipeline Logs.

Completed Execution Logs

It is the end-to-end execution of a pipeline, that is, the record of what happened inside it, from the moment the input passed through the trigger until it ran through the last component in the pipeline.

It is defined by the retention time or amount of messages, whichever comes first. The number of messages are calculated per number of deployed RTUs for a given pipeline.

Pipeline Logs

It is detailed data that allows the user to see in depth what is happening with each execution during the pipeline's running cycle. The limit is given by the retention time or the amount of

stored bytes, whichever comes first. The amount of stored bytes are calculated per number of deployed RTUs for a given pipeline.

Digibee Storage

The Digibee Platform provides a Cloud Storage system for pipelines to read and write files to. The way these files are accessed is through the use of the Digibee Storage Connector. The platform implies limits on the number of bytes stored into this Cloud Storage system.

Object Store

The Digibee Platform provides access to a temporary object storage system that can store any type of JSON object. These objects can be queried based on specific rules.

Production Object Stores are provided in capacity ranges that grow as the number of RTUs grow. They are dedicated instances for every Realm.

Typical Object Store tiers and their match against number of RTUs:

Tier	RTUs	vCPU	RAM
1	Up to 60	0.5	1.7 GB
2	Up to 120	1	3.75 GB
3	Up to 240	2	8 GB
4	Up to 480	4	16 GB
5	Up to 960	8	32 GB

Although the Digibee Capacity team always reviews these tiers so pipelines can make the best use of the Object Stores, customers must design their pipelines in a way that follows Object Stores best practices.

Object Stores computing capacity are not guaranteed in the Test environment, as they are not meant to be heavily used. In this case, they should be used to validate the functional aspects of the pipelines.

Relationship Management

The Digibee Platform provides access to an id relationship management system that can store mappings between data in different systems. The limit implied on the relationship management system is related to the amount of unique mappings that are stored.

Connectivity

The Digibee Platform provides access to our customers' through VPN (Virtual Private Network) and ZTNA (Zero Trust Network Access) connections. Both of them serve the purpose of securing access to resources and granting access to systems, services, and applications based on defined policies.

VPNs and ZTNA

The Digibee Platform provides access to our customers' on-prem systems through private connections. The gateways are sized proportionally according to the number of RTUs in the existing subscription.

Technical Usage Limits

Test Environment Usage Limits

The following limits apply to test RTUs under the current customer's subscription.

Limit	Category	Limit	Scope
Processing capacity	Proportional	64MB / 20% CPU / 10 concurrent executions	Deployed
Egress traffic	Proportional	100 MB per RTU (monthly)	Contracted
Message rate	Proportional	0.1 EPS per RTU	Contracted
Queueing	Proportional	250 Mb per RTU	Deployed
Completed Execution Logs	Proportional	7 days or 604,800 records per RTU	Deployed
Pipeline Logs	Proportional	3 days or 250 MB per RTU	Deployed
Digibee Storage	Proportional	10 MB per RTU	Contracted
Object Store	Range	See object store definition for more details	Contracted
Relationship Management	Proportional	100,000 objects (base) and 10,000 objects (per RTU)	Contracted

Production Environment Usage Limits

Limit	Category	Limit	Scope
Processing capacity	Proportional	64MB / 20% CPU / 10 concurrent executions	Deployed
Egress traffic	Proportional	1 GB per RTU (monthly)	Contracted
Message rate	Proportional	1 EPS per RTU	Contracted
Queueing	Proportional	2.5 Gb per RTU	Deployed
Completed Execution Logs	Proportional	30 days or 2,592,000 records per RTU	Deployed
Pipeline Logs	Proportional	10 days or 1 GB per RTU	Deployed
Digibee Storage	Proportional	100 MB per RTU	Contracted
Object Store	Range	See object store definition for more details	Contracted
Relationship Management	Proportional	10,000,000 objects (base) 1,000,000 objects (per RTU)	Contracted

Usage Limits for All Environments Combined

These limits are not applicable to each environment (Test and Production), they are applicable to your entire realm on Digibee Integration Platform.

Limit	Category	Limit	Scope
VPNs	Proportional	1 GB per RTU (monthly)	Contracted
ZTNA	Proportional	1 GB per RTU (monthly)	Contracted