



PLANNING HIGHER ORG LIMITS FOR SALESFORCE

Abstract

As businesses prosper, they need to cater to a growing customer base with a corresponding increase in the volume of transactions. To do this effectively, SaaS applications must scale to handle the higher load. It is critical for organizations to have diligent planning in place, particularly for multi-tenant platforms like Salesforce so they can predict potential increase in volumes year-on-year (YoY) and for special events such as product launches and promotions to ensure business as usual.

This paper articulates the methodology that can be adopted to calculate the increased org limits for a Salesforce application. To assess this, it is important to understand current usage patterns for critical org limits. Usage patterns must be gathered for an extended period to understand the average and peak utilization patterns for an organization. With this information, projected order volumes and transactions will need to be assessed by working with business units to extrapolate and calculate the necessary increase in limits.

ABSTRACT	1
INTRODUCTION TO LIMITS	2
DEFAULT ENTITLEMENTS AND CRITICAL LIMITS	3
APPROACH TO MANAGE HIGHER ORG LIMITS	4
UNDERSTANDING CURRENT ORG LIMITS AND USAGE	4
PLANNING FOR INCREASED ORG LIMITS FOR HIGHER USAGE	7
STRATEGIES TO OPTIMIZE LIMIT CONSUMPTION	8
AUTOMATED MONITORING	8
CONCLUSION	10
REFERENCES	11
GLOSSARY	11
ABOUT THE AUTHOR	12
ACKNOWLEDGEMENTS	12

Introduction to Limits

Limits on multi-tenant platforms are designed to protect shared resources. To ensure that the Salesforce platform is available to all customers and to maintain optimum performance, transaction loads are balanced by Hyperforce IaaS by imposing various engineering base-limits on each of the customer orgs. Contractual limits based on a customer's contract may also determine the allocations for that customer. All limits have a default allocation, which is enforced daily to ensure fair sharing of resources in a multi-tenant environment and to protect the service.

Some of the governor limits are called soft limits and higher entitlements for soft limits can be provisioned through purchase or by acquiring additional licenses. Few examples of soft limits include Async Apex Execution limits, HVPE add-on, workflow time trigger limit per hour, and data storage. Some of these changes require a larger turnaround and review with Salesforce engineering teams, whereas some of the limits can be raised temporarily through a Salesforce support case raised a couple of weeks in advance. An example of a soft limit that can be increased through a Salesforce case is Max Org-Wide Long-Running Apex Threads which has a default value of 10. This limit pertains to the number of synchronous concurrent long-running transactions that take 5 seconds or more for each org.

Hard limits are mostly related to apex programming limits and other org limits which cannot be increased through purchase. Consider a Salesforce org which has implemented Salesforce flow to enrich the contact record through a flow every time a new contact is created. This will consume the Paused and Waiting Flow Interviews limit which is a hard limit of 50,000 per org and when the number of records in the flow table reaches this limit, an unhandled system error is raised. Another example of a hard limit is the Flex Queue jobs which has a limit of 100. This means at any



point of time a maximum of 100 batch jobs can be placed in Hold status for future execution.

Default Entitlements and Critical Limits

Some of the default entitlements for a customer can be viewed through the Setup – Company Information and Usage Based Entitlements screens. The Lightning Platform REST API also provides details of the maximum allocation and the remaining allocation at the current point of time.

Syntax : /services/data/vXX.X/limits/

Identification of the critical limits for a particular customer depends entirely on the architecture landscape, implementation, and the integration patterns deployed.

For example, if a customer has implemented Salesforce Connect to integrate external data sources, then the OData Hourly limit which has a default value of 20,000 per hour, needs to be analyzed. If a customer utilizes Industries Order Management (IOM) for order orchestration and fulfillment, then the Asynchronous Apex method executions and Event Delivery limits need to be analyzed. There are other means by which these limits can be consumed, so careful evaluation of the implementation is necessary.

In this paper, we will refer to a few common critical org limits such as Asynchronous Apex method executions, Event Delivery, Event publishing, and data storage. Table 1 illustrates a sample way to list all the critical limits for a customer.

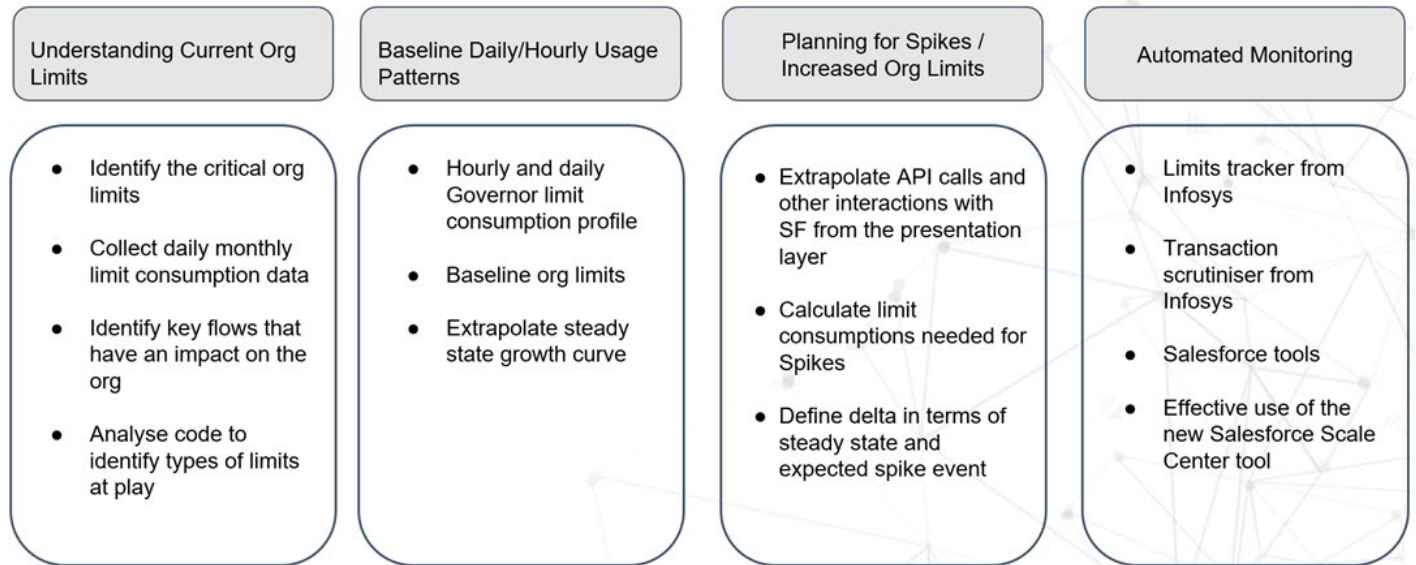


Limit	Hourly	Daily	Monthly	Comments
Event Delivery		550,000 ¹	1,550,000 ¹	Refers to the maximum number of event notifications delivered in the last 24 hours - can be delivered to CometD and Pub/Sub API subscribers. Limit: 50,000. Customers can buy the HPVE Add-on which will provide an additional 500,000 limits.
Event Publishing	250,000 ¹			Refers to the maximum number of event notifications that are published per hour. 250,000 included with org license + 25,000 purchased from add-on license (optional).
Asynchronous Apex method executions		250,000 ¹		Daily number of Asynchronous Apex method executions, which includes batch Apex, future methods, queueable Apex, and scheduled Apex. Limit - 250,000 or 200 multiplied by the number of user licenses in the org (whichever is greater).
Data Storage				10GB + (120MB x number of licenses)

Table 1 – Template listing critical limits for a customer.

¹Indicative as of 2023 Enterprise, Performance, Unlimited editions. Reference section ‘Salesforce Governor Limits Ref’

Approach To Manage Higher Org Limits



Understanding Current Org Limits and Usage

There are a couple of strategies to evaluate the current usage patterns of a customer.

1

Calculate org limits that are critical for each transaction/process flow

2

Review the daily/hourly usage pattern for key org limits

Calculate org limits that are critical for each transaction/process flow

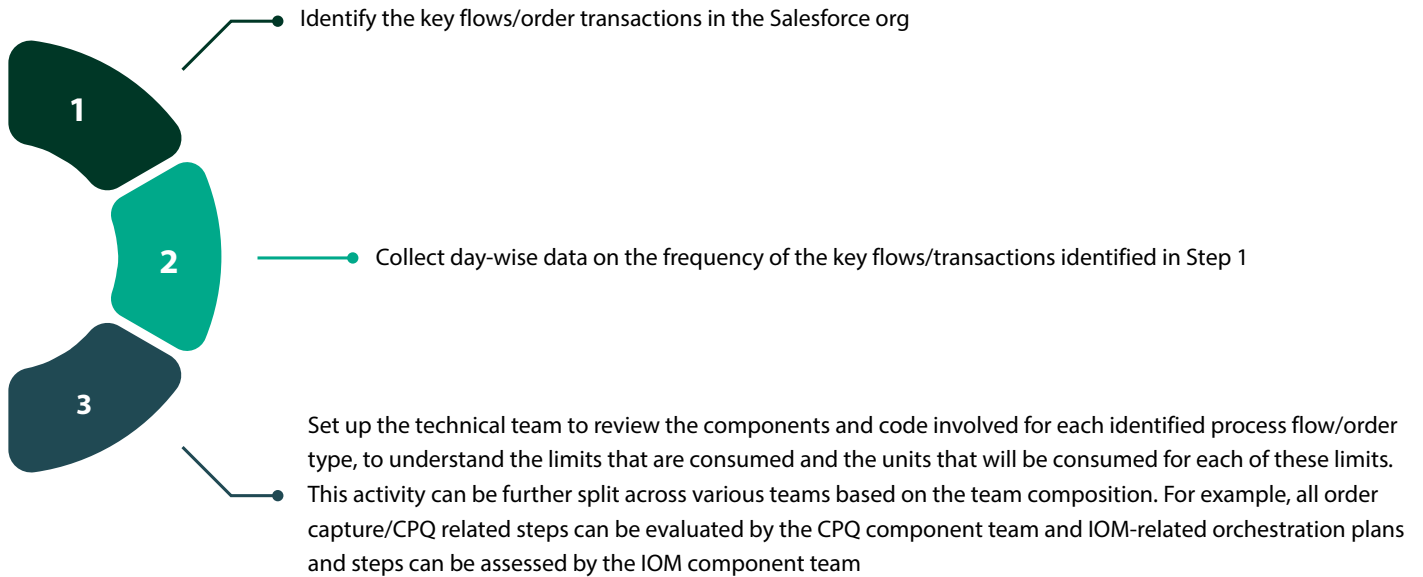


Table 2 illustrates a sample template that can be utilised to collect monthly data for the Async Apex Execution limit.

Transaction Type	1/5/2023	2/5/2023	3/5/2023	4/5/2023	6/5/2023	7/5/2023	8/5/2023
SIM – Activation	15	20	10	20	20	13	18
SIM – Exchange	39	45	45	41	39	40	42
Plan Change	121	130	76	98	80	100	125
Data Update	3				10	1	
Disconnect	145	158	214	186	115	134	103
Modify	1645	1882	1912	780	2000	890	1391
New	36	42	43	50	38	52	77
Port-In	38	34	16	22	25	24	40
Re-Contract	15	29	29	21	20	26	40
Reconnect	1226	72	1090	358	1400	381	103
Suspend	2261	10	1828	4	2420	735	4
Grand Total	5544	2422	5263	1580	7130	2396	1943

Table 2 – Sample template to collect monthly data

Choose the limits for a date which has an acceptable daily BAU peak for computation purposes. For this example, data from 6/5/2023 is used to compute the Async Apex limit for various transactions in the system. In the illustration in Table 3, 87580 Async Apex Execution limits were consumed by the org on 6/5/2023.

Type	Number	Unit Consumption of Async Apex Execution*	Async Apex Execution
SIM – Activation	20	30	600
SIM – Exchange	39	30	1170
Plan Change	80	30	2400
Data Update	10	30	300
Disconnect	115	30	3450
Modify	2000	20	40000
New	38	20	760
Port-In	25	20	500
Re-Contract	20	20	200
Reconnect	1400	10	14000
Suspend	2420	10	24200
		Total	87580

Table 3 – Illustration of apex async execution limits consumed by the org on a specific date

The first row in Table 3 signifies that there were 20 transactions of type 'SIM – Activation' on 6/5/2023 and each of those transactions consumed 30 units of Async Apex Execution limits totaling to 600 units.

* The unit consumption can be different for each type of transaction and will need to be evaluated by the technical team based on the configuration and customization done.

Review the daily/hourly usage pattern for key org limits

Raise a Salesforce case or work with the Technical Account Manager to get the hourly or daily limits consumed (based on what type of limit it is) across a whole month or an extended period as agreed with business stakeholders. Identify the peak usage for all critical limits. Use this as a baseline to make projections for higher volumes in future. The data sets can be collated for a couple of months, if required, to make projections on the percentage increase.

Table 4 illustrates the daily Platform Events Delivered limit for each platform event published in the Salesforce platform. In Table 4, numDelivered denotes the number of platform events delivered and ChannelName denotes the name of the platform event. This snapshot will be provided by Salesforce on request. In this illustration, 10,510 Platform Events Delivered limits were consumed on 30/10/2022.

Time	Num Delivered	Channel Name	Topic	Platform Type
2022-10-30T00:00:00.000	8414	SharedSubscriber /event/01I2w000000IB5B	/event/01I2w000000IB5B	High Volume
2022-10-30T00:00:00.000	1420	/eventAccountHierarchyDetails__e	/event/01I2w000001jZJP	High Volume
2022-10-30T00:00:00.000	11	/event/FileEvent__e	/event/01I2w000001jQ0n	High Volume
2022-10-30T00:00:00.000	85	/event/File_Message_Event__e	/event/01I2w000000JqVv	High Volume
2022-10-30T00:00:00.000	580	/event/3rd PartyPlatformEvent__e	/event/01I2w000001jMV3	High Volume
Total	10,510			

Table 4 – Sample of daily Platform Events Delivered limit for each platform event published in the Salesforce platform

Planning for Increased Org Limits for Higher Usage

Planning for BAU

During this phase, data on the number of process flows/transactions/orders executed each day in production is used to calculate the consumption limit per day for each critical limit identified for the customer. This information can now be augmented with the projected expected daily volumes of transactions for the next 1 year, 2 years, and so on in consultation with the business teams.

Channel	Journey/order type/ process flow	Current daily volumes in production org	Expected daily volume from a 1-year perspective	Expected daily volume from a 2-year perspective
Web	Plan Change	40	60	75
	Port-In	10	20	25
	Terminate	5	10	15
Retail	Plan Change	60	80	100
	Terminate	10	15	20
App	Subscribe VAS	50	75	100
	Unsubscribe VAS	10	12	15

Note: All numbers are for illustration purposes only

Projected increased limits = Limit consumption per day x expected daily volumes from a 1-year perspective

Customers undergoing a multi-year digital transformation must consider the order/transaction volumes in both the legacy CRM and in the Salesforce target stack to estimate the total increase in org limits required.

Planning for Special Events

Customers typically have key events once or twice yearly which may lead to a sudden surge in the transactions on the platform. To successfully run these events, customers must ensure that the platform is capable of handling the additional traffic and that BAU transactions also run smoothly during this period.

Examples of special events from various industries:

- » In the financial services industry, insurance clients see a surge in policy applications in the last month of the financial year for availing income tax benefits
- » In the retail industry, it is common for high volumes during special events such as end of the year or Independence Day sales
- » In the telecom industry, the launch of the latest iPhone, Samsung, and other brands' flagship models are considered key events of customer interest

Table 5 and Table 6 provide sample illustrations of how org limits can be calculated for special events.

# of iPhone/Samsung pre-orders expected per day	2000
# of Async Apex Execution Limit consumed per order	8
Total Async Apex Execution Limit required	16000

Table 5 – Illustration to calculate org limit for iPhone/Samsung pre-order

# of GST/VAT/Taxation change orders submitted in a day	350000
# of Async Apex Execution Limit consumed per order	4
Total Async Apex Execution Limit required	1400000

Table 6 – Illustration to calculate org limit for GST/VAT/Tax price update on existing assets



Strategies to Optimize Limit Consumption

- **Optimizing the implementation**
In Industries Order Management, each step in the orchestration plan which is of type Auto Task or Callout consumes an Async Apex Execution limit and a Platform Event Published limit. Designers should consider reducing the number of Auto Tasks in the plan and perform the same actions with a fewer number of Auto Tasks, wherever feasible. Along similar lines, instead of making 3 or 4 callouts to an external system, evaluate design patterns where a single payload with all the information can be sent to the external system to process. The same rationale can be applied to Salesforce Apex code which implements future or queueable classes which consume the Async Apex Execution limit.
- **Controlling Load Using an External Framework**
In scenarios where external systems are creating bulk orders in Salesforce or triggering Salesforce APIs, incoming requests can be throttled by the external system, by reducing the number of parallel threads used to invoke Salesforce transactions, to stay within permissible limits.

Automated Monitoring Infosys Automated Monitoring Tools


Infosys code store has a few automated monitoring tools to help monitor the Salesforce limits:

- **Limit Tracker for Salesforce**
 - o Limit Tracker tracks the consumption of various platform limits in Salesforce and provides notification capability
- **Transaction Scrutinizer**
 - o Transaction Scrutinizer calculates the percentage consumption of different governance in a transaction such as CPU limit, DML rows, DML statement, and SOQL rows and has the capability to send automated emails to users based on the defined threshold. For example, if CPU limit consumed is greater than 70% of the available limit then an email notification is sent to the administrator.

Salesforce Monitoring Tools

• Scale Center

- o Scale Center is a Salesforce product launched in June 2023 that enables customers to proactively identify and assess performance and scale hotspots by providing near real-time insights. It can be enabled by searching for Scale Center using the quick find box and turning on Enable Scale Center
- o Performance analysis reports can be generated for defined periods to compare org performance metrics.
- o The following metrics will be part of the consolidated report generated:



Average Request Time (ms) ⓘ

Average time for the platform to process and send requests.

Click and drag on any chart to select a time range for analysis.




Figure 1 – Consolidated performance analysis report (Sample Illustration only for reference from Salesforce Scale Center)

Conclusion

Customers and partners can leverage the strategies outlined in this paper to assess their future scalability needs related to org limits for their Salesforce implementations. This assessment is critical as customers need to strike a balance between reducing their opex and ensuring seamless BAU with the required increase of org limits.

- Org limit calculations and projections must be revised based on new business functionality delivered with each release and actual transaction volumes observed in the production org
- We strongly recommend real-time monitoring of org limits with email alerts for breaches in the thresholds set for each critical limit
- For any special events such as an iPhone launch for telecom operators, customers will need to work with the Salesforce account executives to set increased org limits which may need to be provisioned. Customers who have signed up for Salesforce Signature Success may be covered for one special event every year
- Designers and developers must keep tabs on the various limits that each solution will consume starting from the analysis phase of the delivery cycle, to prevent issues pertaining to limit breaches at a later stage. Knowledge of the expected number of transactions/volumetrics at the discovery stage is critical to help evaluate the impact on the limits
- It is recommended to keep a buffer on the calculated org limits, to avoid unexpected spikes or unplanned overlap between two or more high volume events. Typically, customers are seen allocating 10% or higher as buffer

Including org limit consumption projections in the early stages of delivery, gathering real-time consumption data for future projections, and revising org limits based on actual transaction volumes will ensure that SaaS applications are well-tuned and can scale rapidly based on evolving business requirements.



References

- Maximum Paused and Waiting Flow Interviews - <https://help.salesforce.com/s/articleView?id=000389239&type=1>
- Limits REST API - https://developer.salesforce.com/docs/atlas.en-us.244.0.api_rest.meta/api_rest/resources_limits.htm
- Process Entitlements - https://help.salesforce.com/s/articleView?id=sf.process_usage_entitlements.htm&type=5
- Apex Flex Queue - https://help.salesforce.com/s/articleView?id=sf.code_apex_flex_queue.htm&type=5
- Platform Event Entitlements - https://developer.salesforce.com/docs/atlas.en-us.platform_events.meta/platform_events/platform_event_limits.htm
- Salesforce Governor Limits Ref - https://developer.salesforce.com/docs/atlas.en-us.salesforce_app_limits_cheatsheet.meta/salesforce_app_limits_cheatsheet/salesforce_app_limits_platform_apexgov.htm - <https://salesforce.stackexchange.com/questions/213322/how-is-calculated-24-hour-period-for-the-limitation-of-asynchronous-apex-execu>
- Salesforce Scale Center - https://help.salesforce.com/s/articleView?id=sf.scale_center_overview.htm&type=5



Glossary

API	Application programming interface
BAU	Business as usual
CRM	Customer relationship management
CPQ	Configure price and quote
GST	Goods and Services tax
IOM	Industries order management
VAT	Value Added tax

Author



Karthik Prasad

Karthik Prasad is a Senior Technology Architect with around two decades of experience in CRM application enterprise architecture and platform advisory services for cloud-based applications. He is also a Salesforce certified Application Architect and a Salesforce Communications Cloud Accredited Professional. As a platform and solution architect he has been part of numerous digital transformation journeys across various industry verticals and specializes in Salesforce Industries implementations.

Acknowledgements

I would like to acknowledge and offer my gratitude to the many people who have contributed to this white paper. Their support has been invaluable.

In particular, I wish to thank my colleague and subject matter expert Dr Thejasvi N, Industry Principal, who helped with the necessary data points, case studies, and shared experiences.

I am also grateful for the reviews and feedback from Kannan N, Senior Principal Architect, Deepak Jha, Industry Principal, Anjali Y, AVP, Arnab Das, Principal Architect (Fluidio) and Sushil Hinduja, Industry Principal.



Infosys Cobalt is a set of services, solutions and platforms for enterprises to accelerate their cloud journey. It offers over 35,000 cloud assets, over 300 industry cloud solution blueprints and a thriving community of cloud business and technology practitioners to drive increased business value. With Infosys Cobalt, regulatory and security compliance, along with technical and financial governance comes baked into every solution delivered.

For more information, contact askus@infosys.com



© 2023 Infosys Limited, Bengaluru, India. All Rights Reserved. Infosys believes the information in this document is accurate as of its publication date; such information is subject to change without notice. Infosys acknowledges the proprietary rights of other companies to the trademarks, product names and such other intellectual property rights mentioned in this document. Except as expressly permitted, neither this documentation nor any part of it may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, printing, photocopying, recording or otherwise, without the prior permission of Infosys Limited and/or any named intellectual property rights holders under this document.