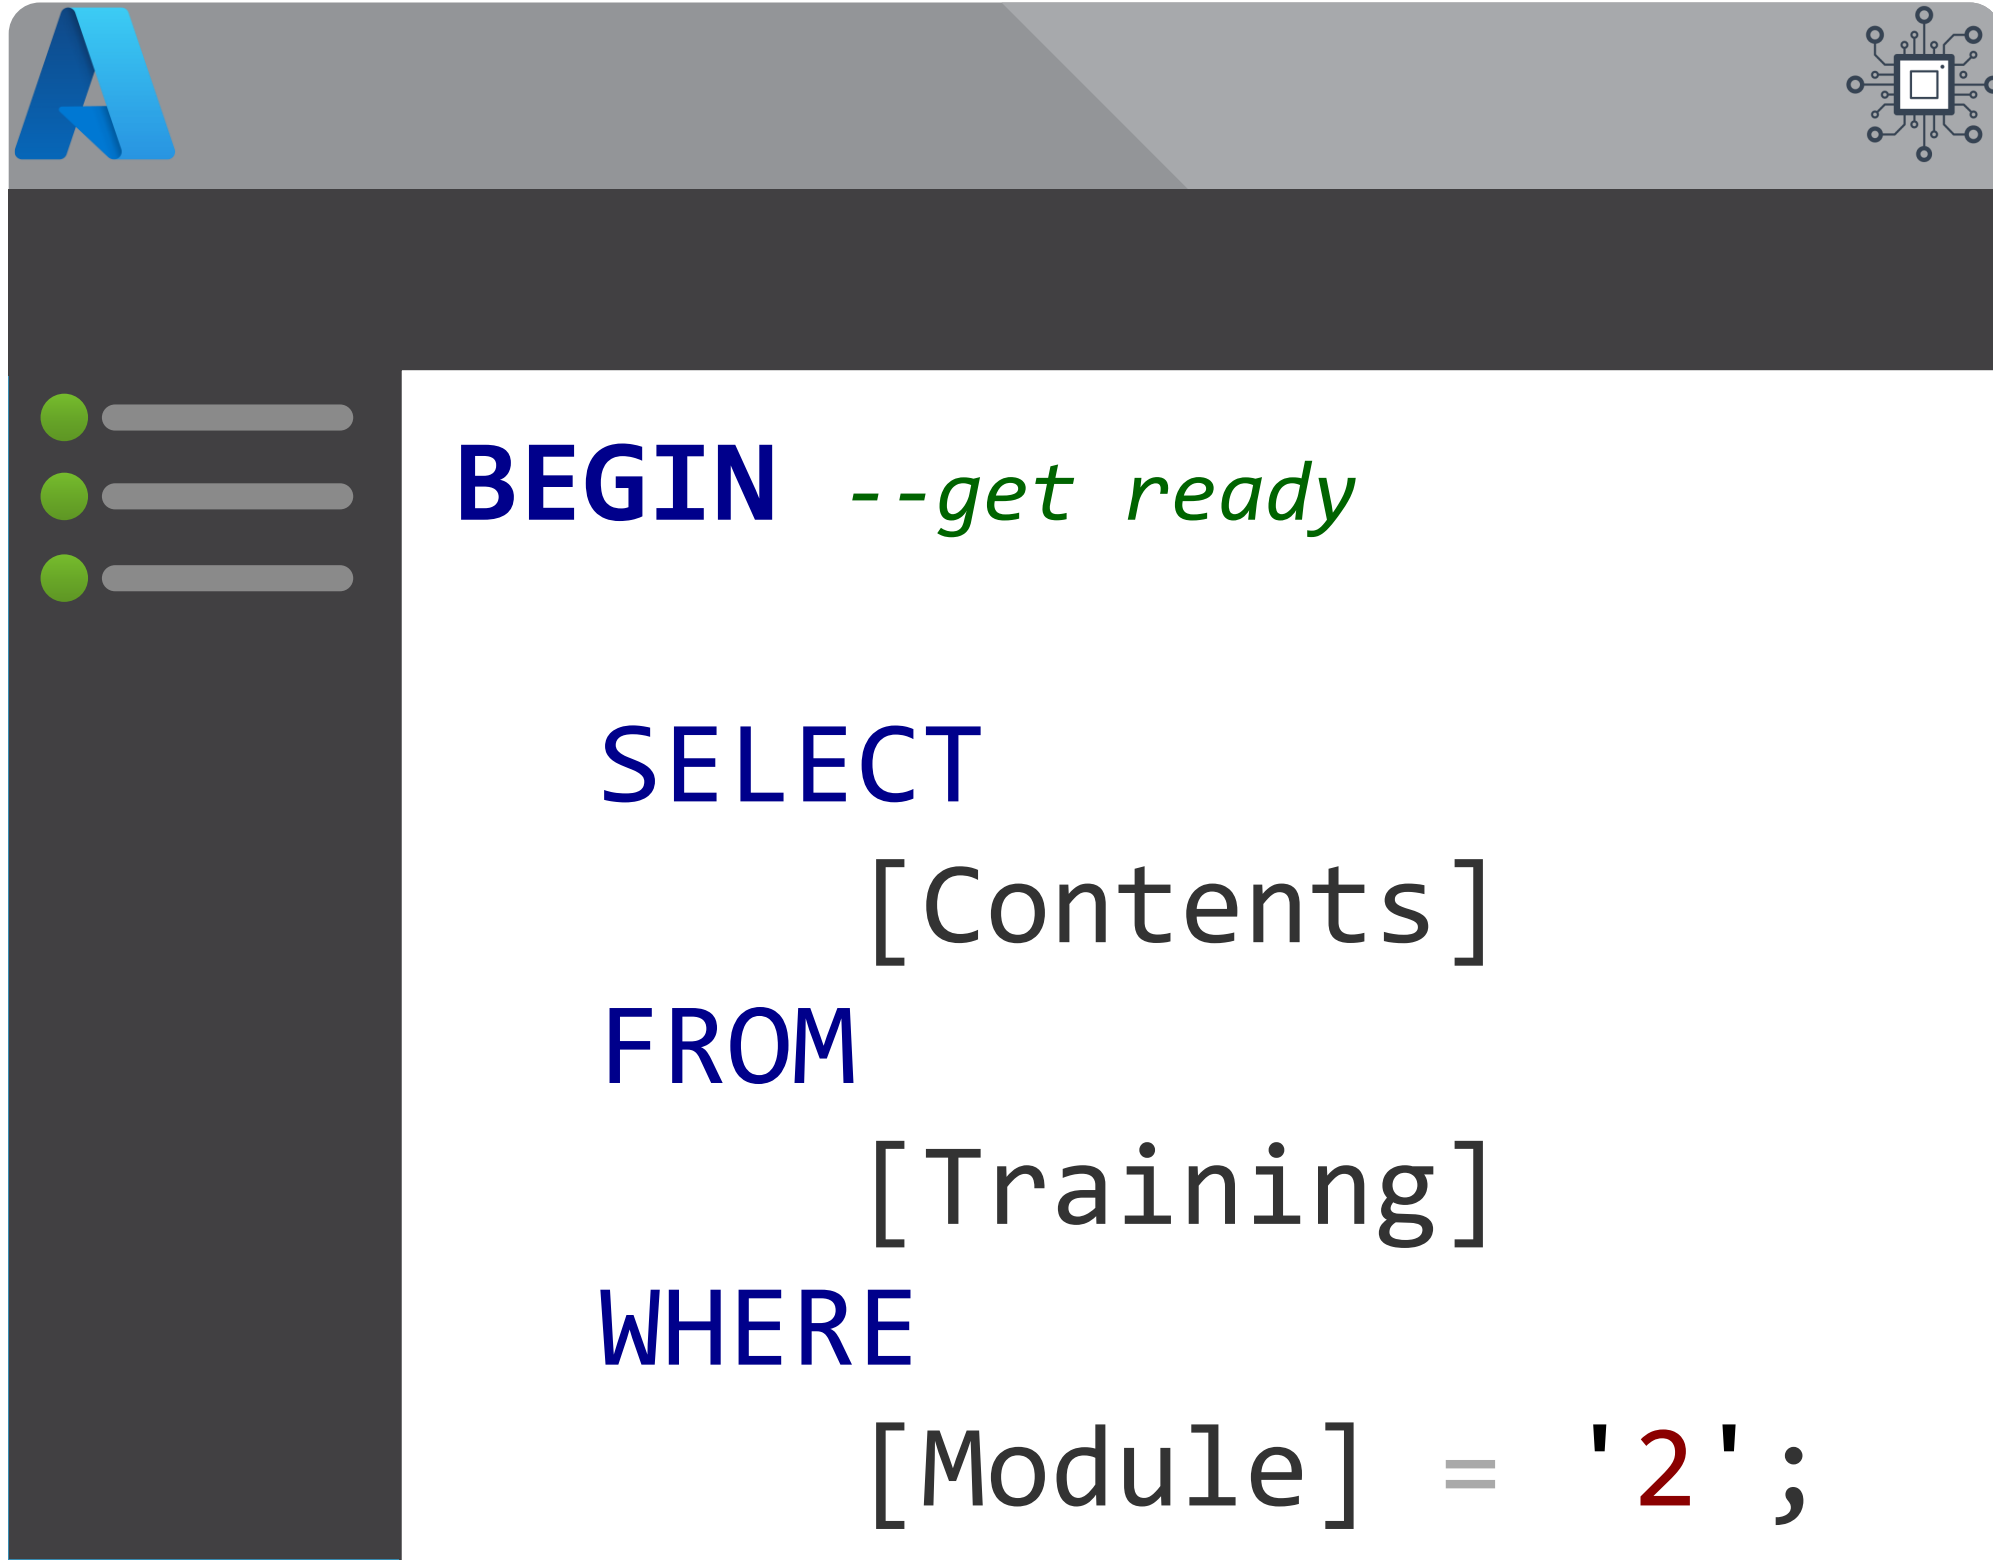


Module 2


Integration Runtime Design Patterns



- Compute Types
 - Azure
 - Hosted
 - SSIS
- Patterns & Configuration

Module 2

Integration Runtime Design Patterns



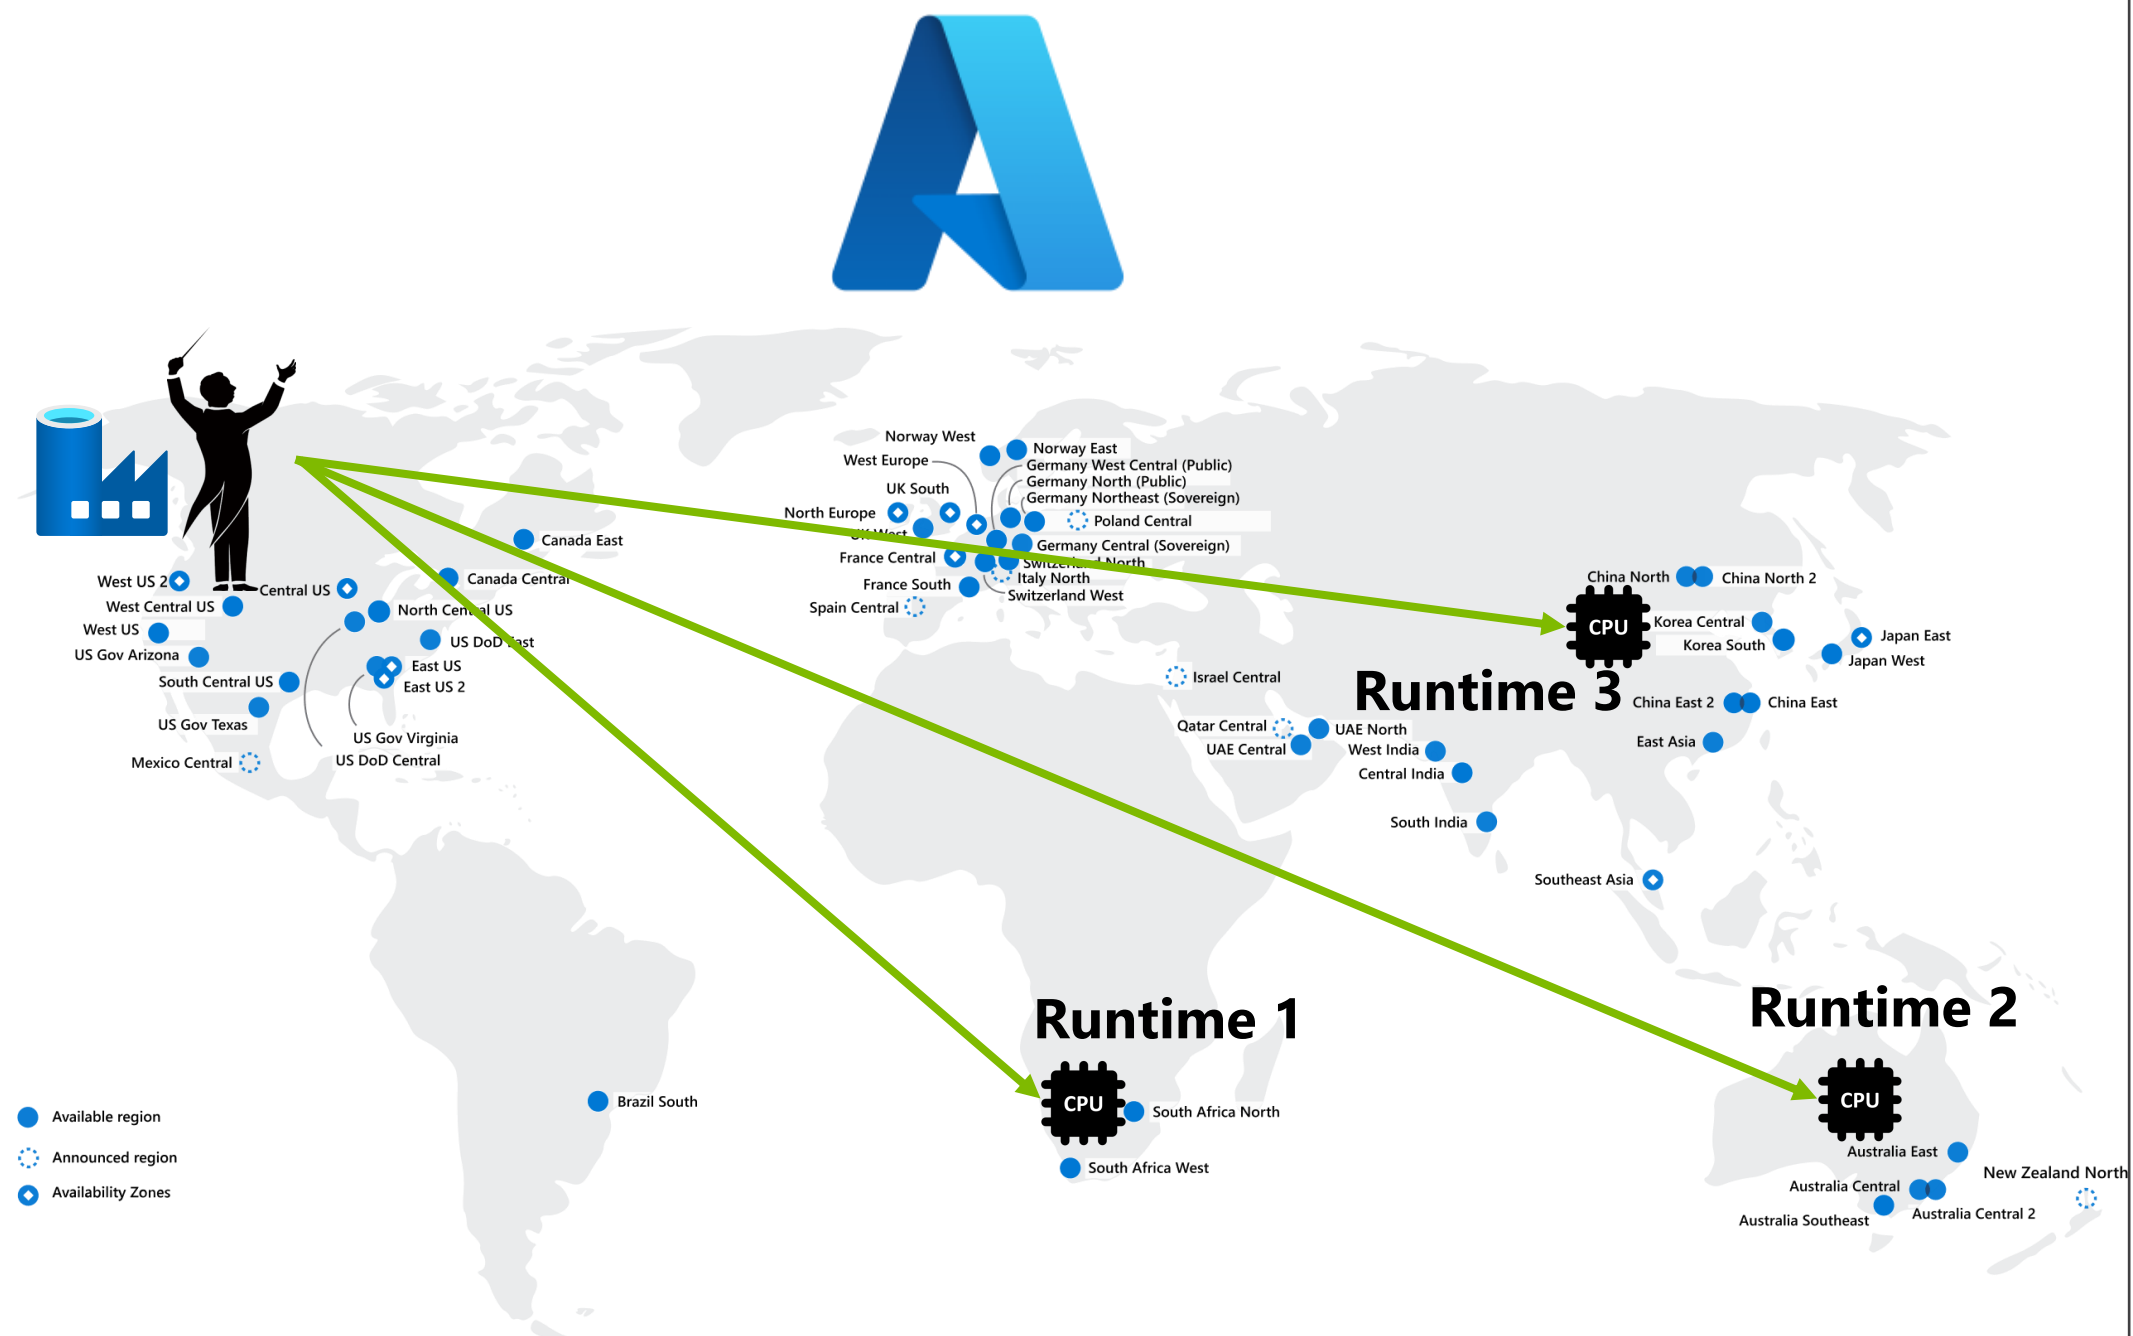

- Compute Types

- Azure
- Hosted
- SSIS

- Patterns & Configuration




What is an Integration Runtime?






What can an Integration Runtime do?



Module 2



Integration Runtime Design Patterns




- Compute Types
 - Azure
 - Hosted
 - SSIS
- Patterns & Configuration





What can an Integration Runtime do?





What is an Integration Runtime?




Orchestrator

A Fixed
Region

Runtime
Flexible
Location



AutoResolveIntegrationRuntime



Edit linked service

Azure SQL Database [Learn more](#)

To avoid publishing immediately to Data Factory, please use Azure Key Vault to retrieve secrets securely. [Learn more](#)

Name *

AzureSqlDatabase1

Description


Connect via integration runtime *

AutoResolveIntegrationRuntime

Connection string

Azure Key Vault

Source
LS



Source

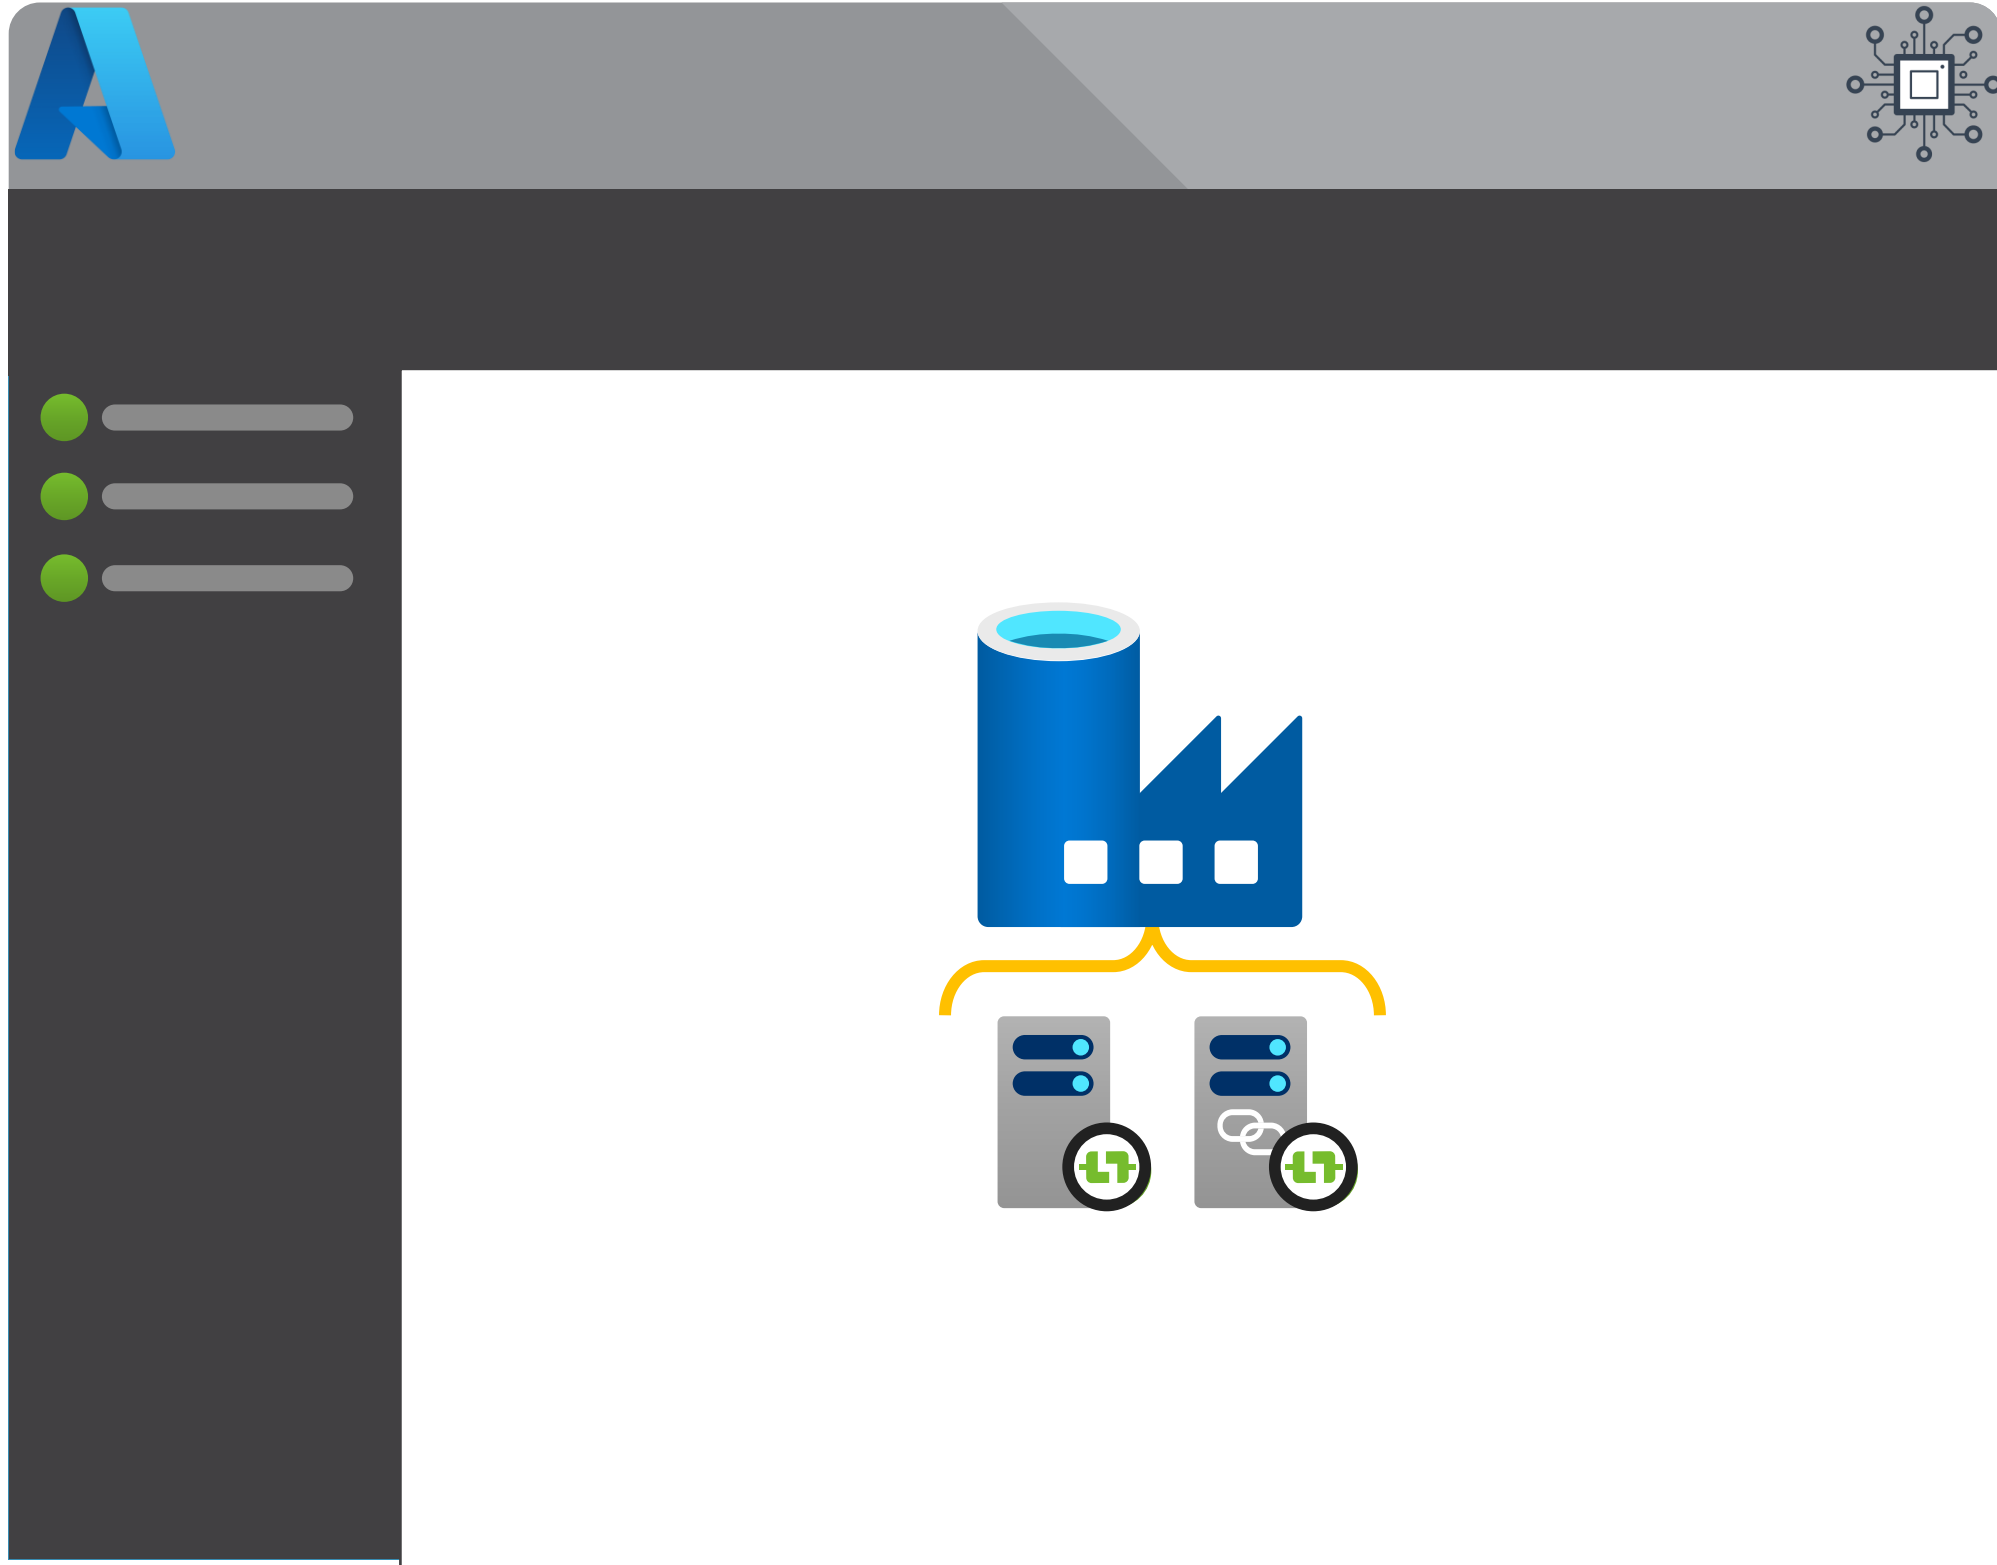
Sink

Sink
LS




Module 2

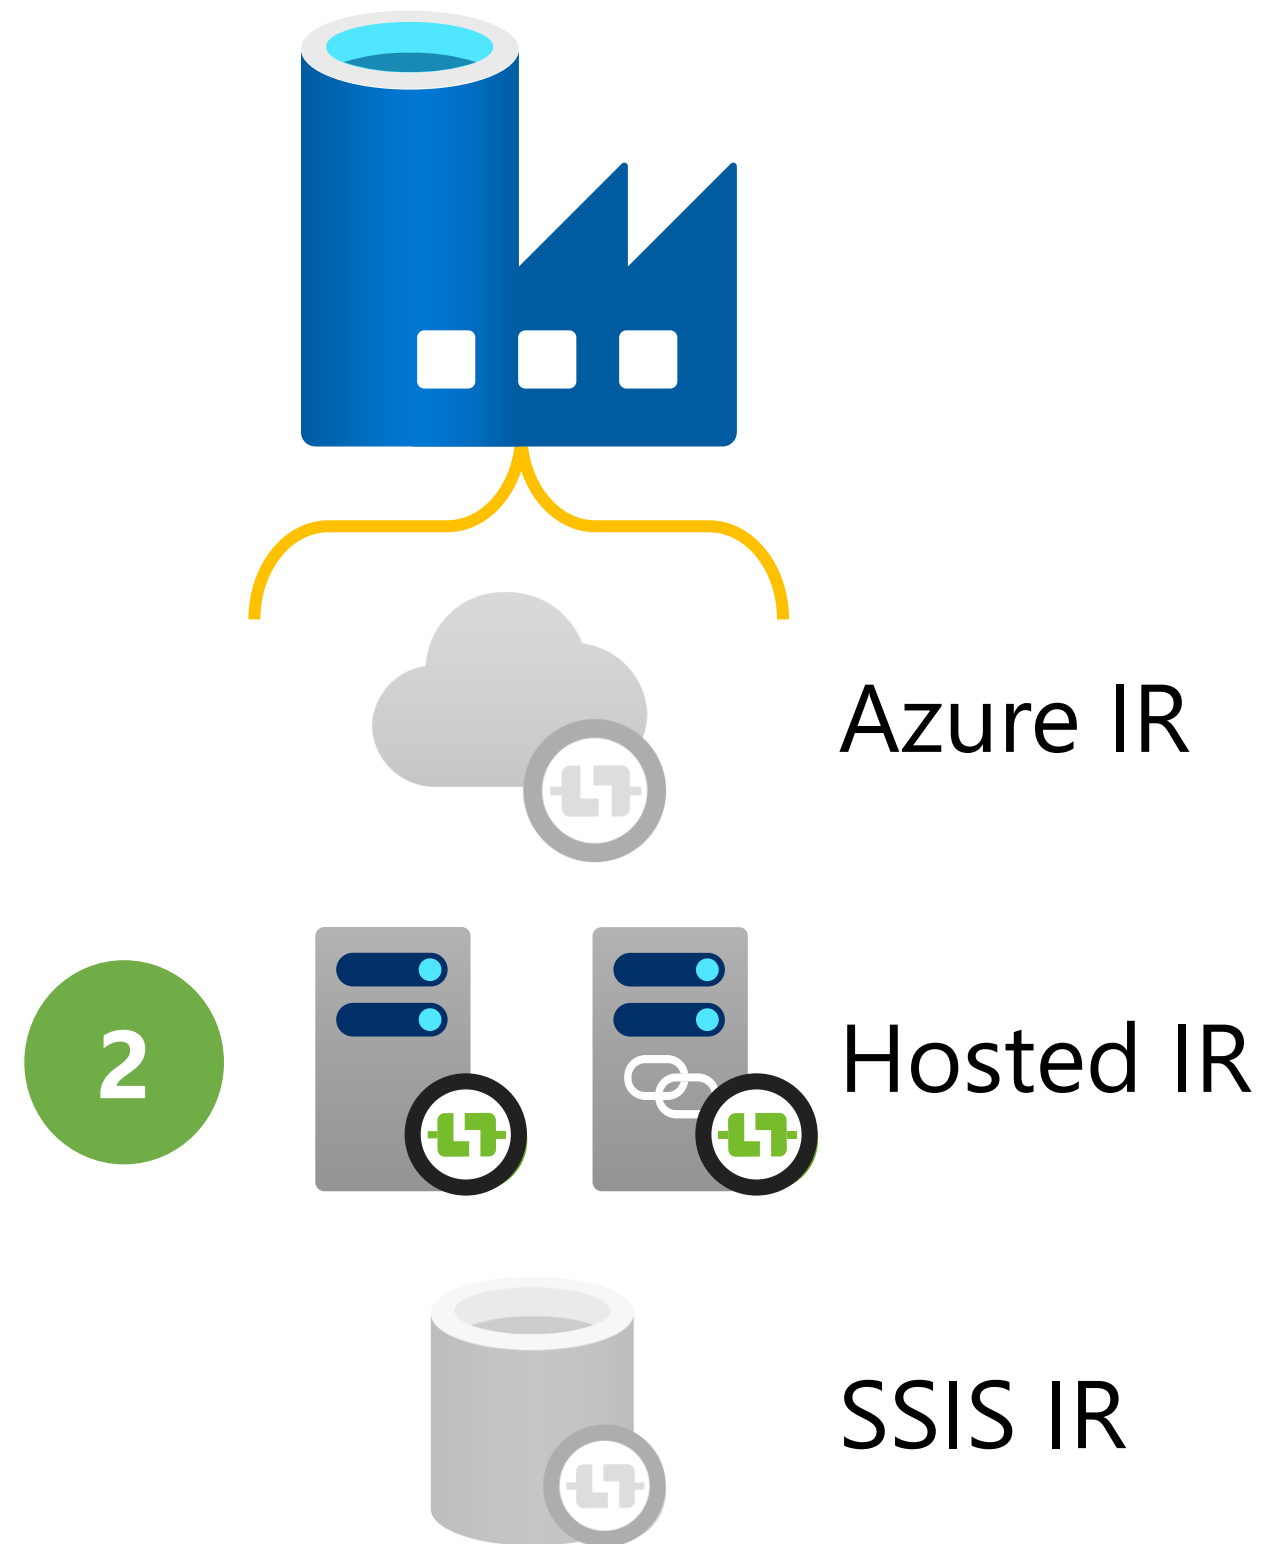

Integration Runtime Design Patterns

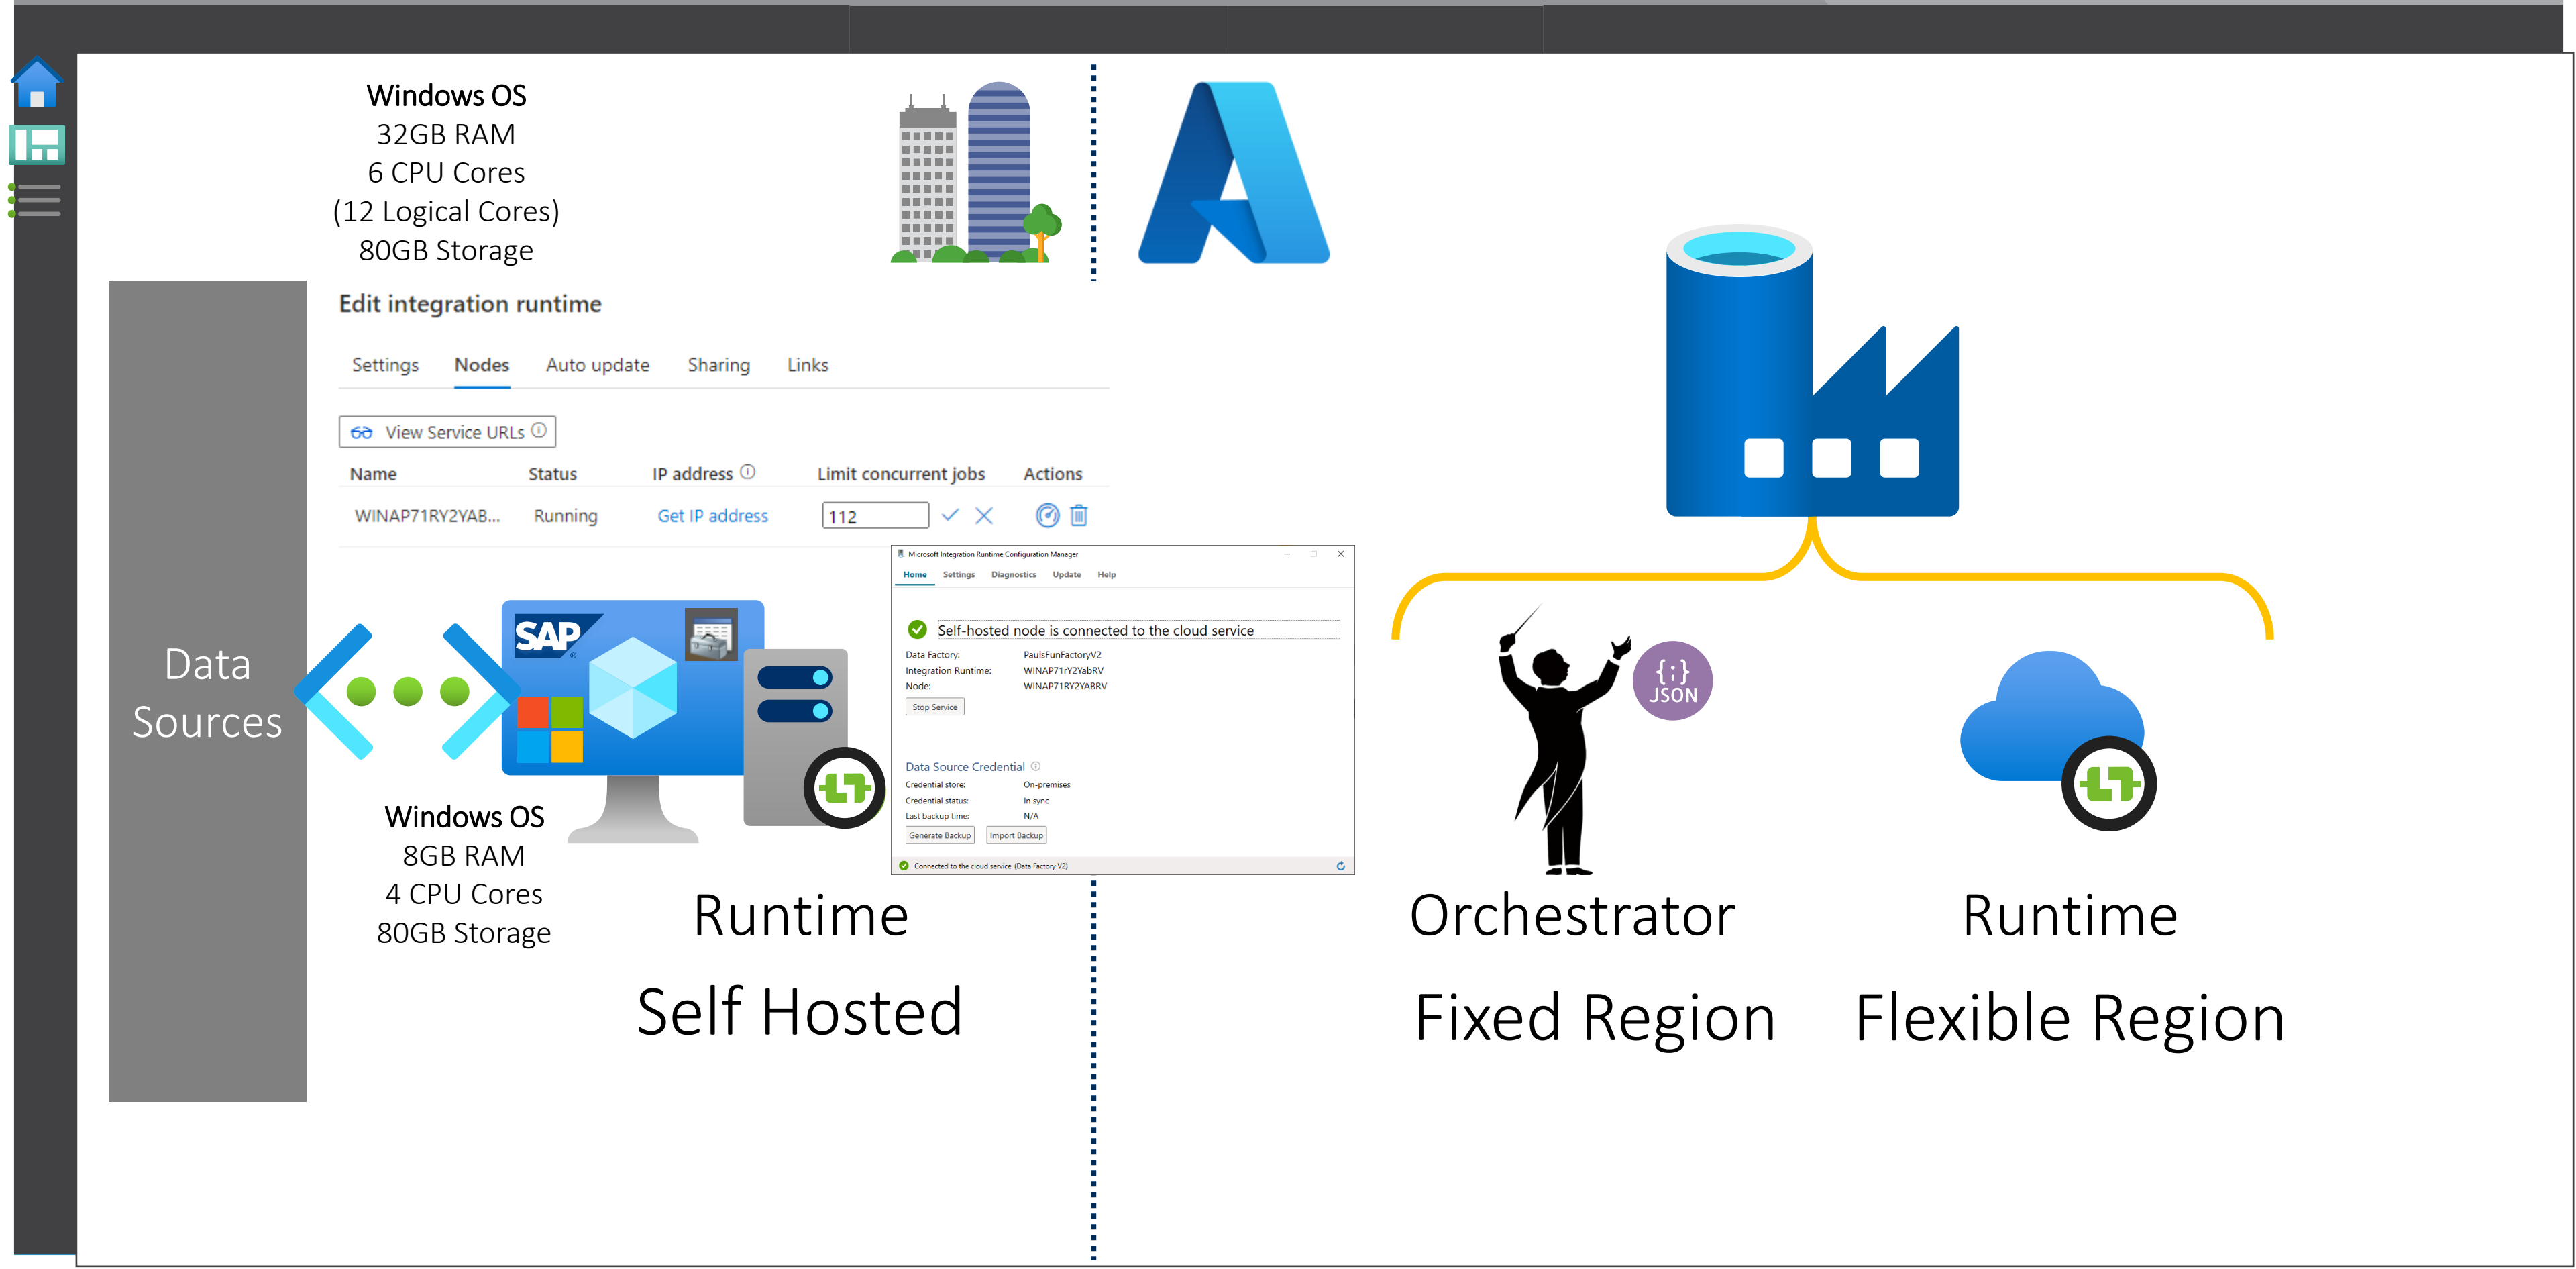


- Compute Types
 - Azure
 - Hosted
 - SSIS
- Patterns & Configuration






What can an Integration Runtime do?








Hosted Integration Runtime



Runtime
Self Hosted



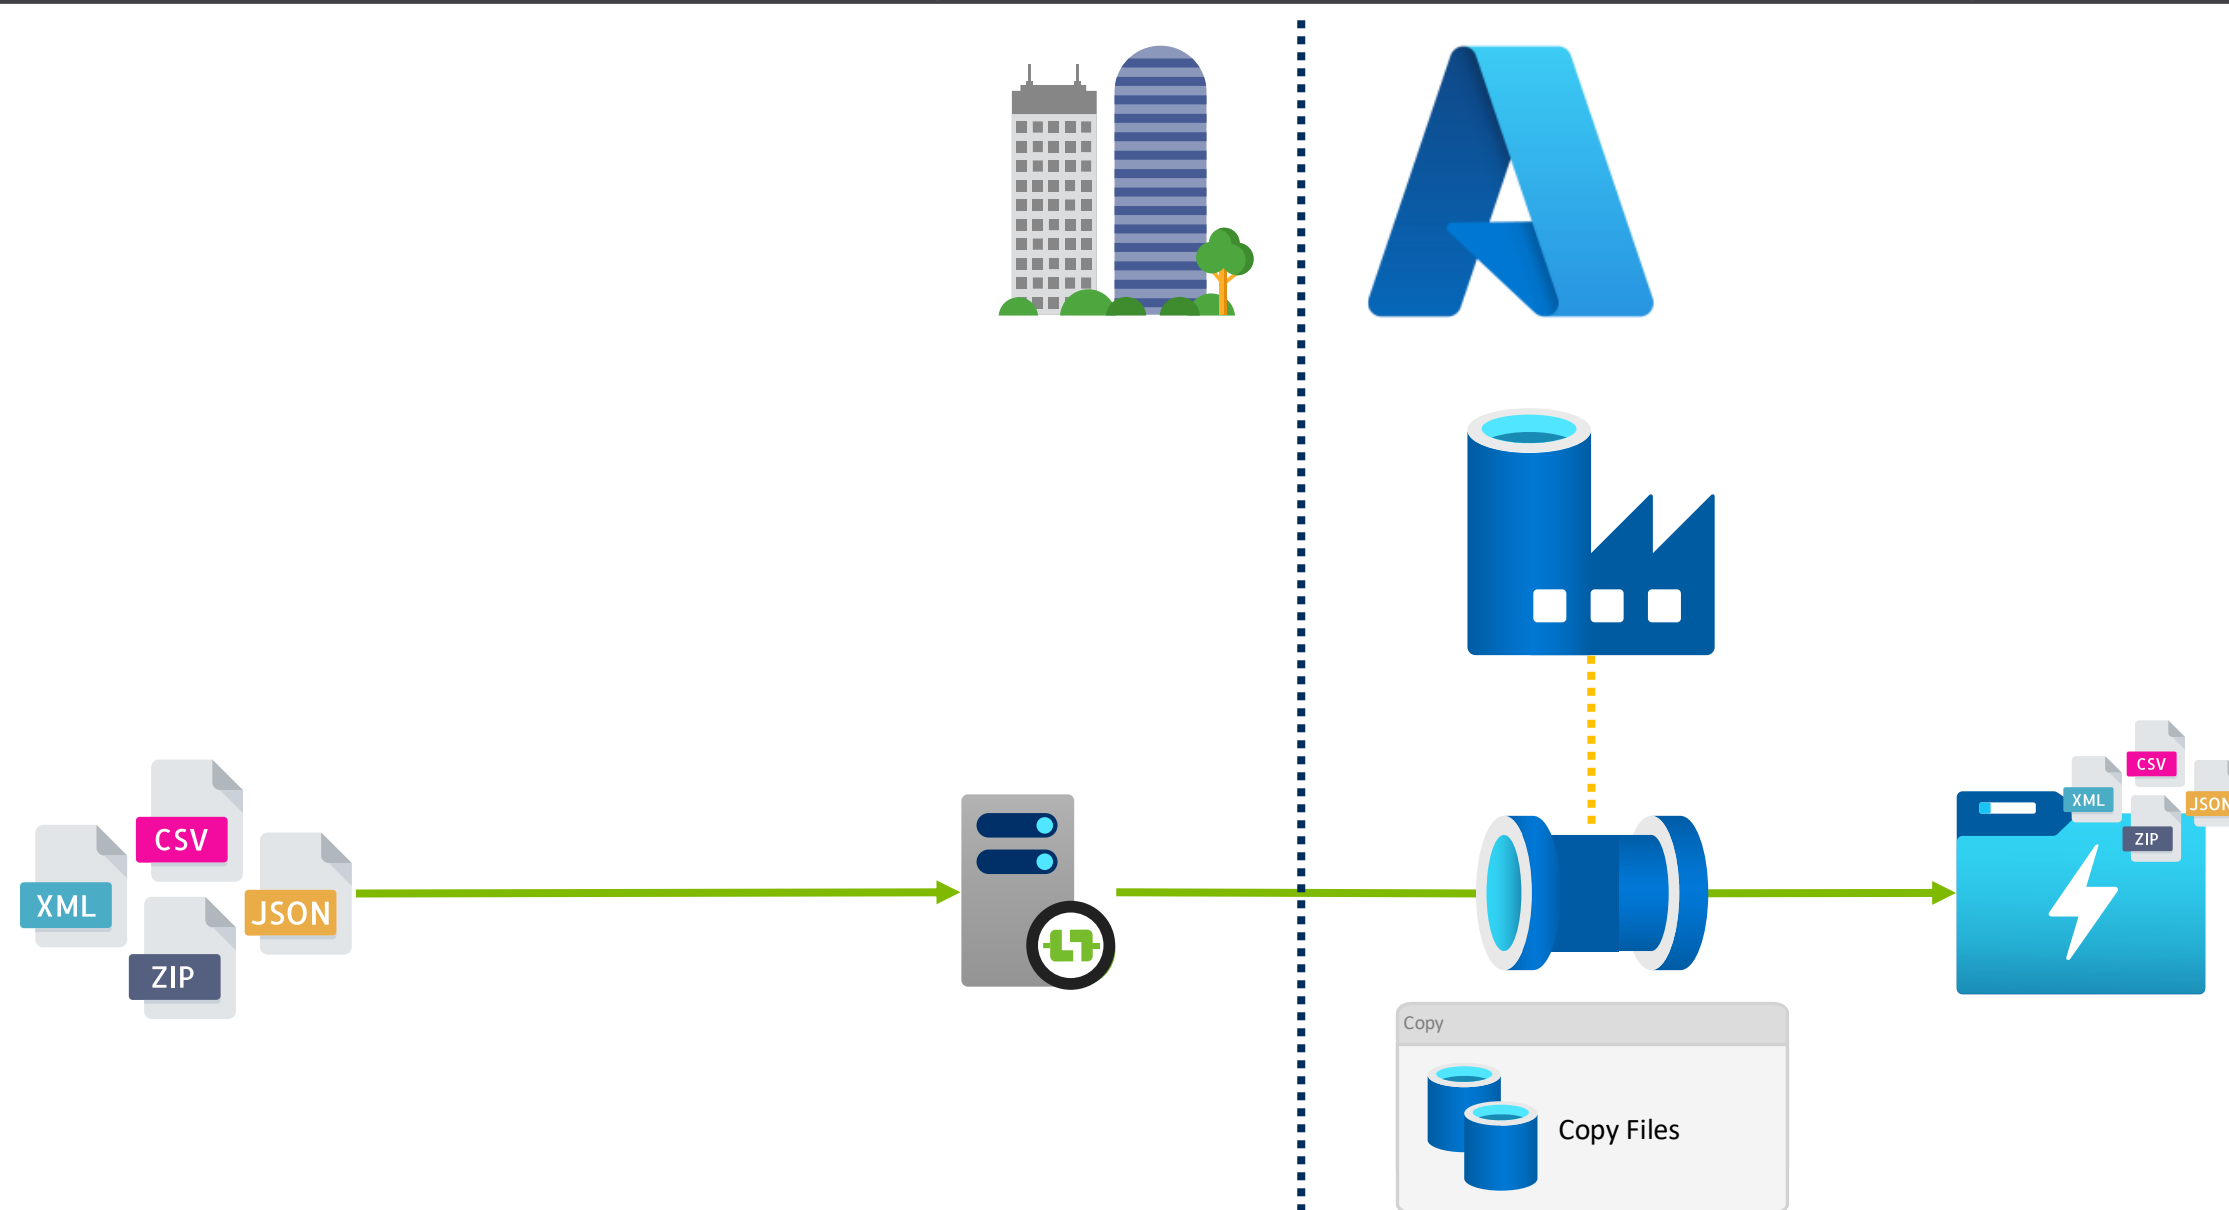
Orchestrator
Fixed Region




Runtime
Flexible Region






Single Hosted IR



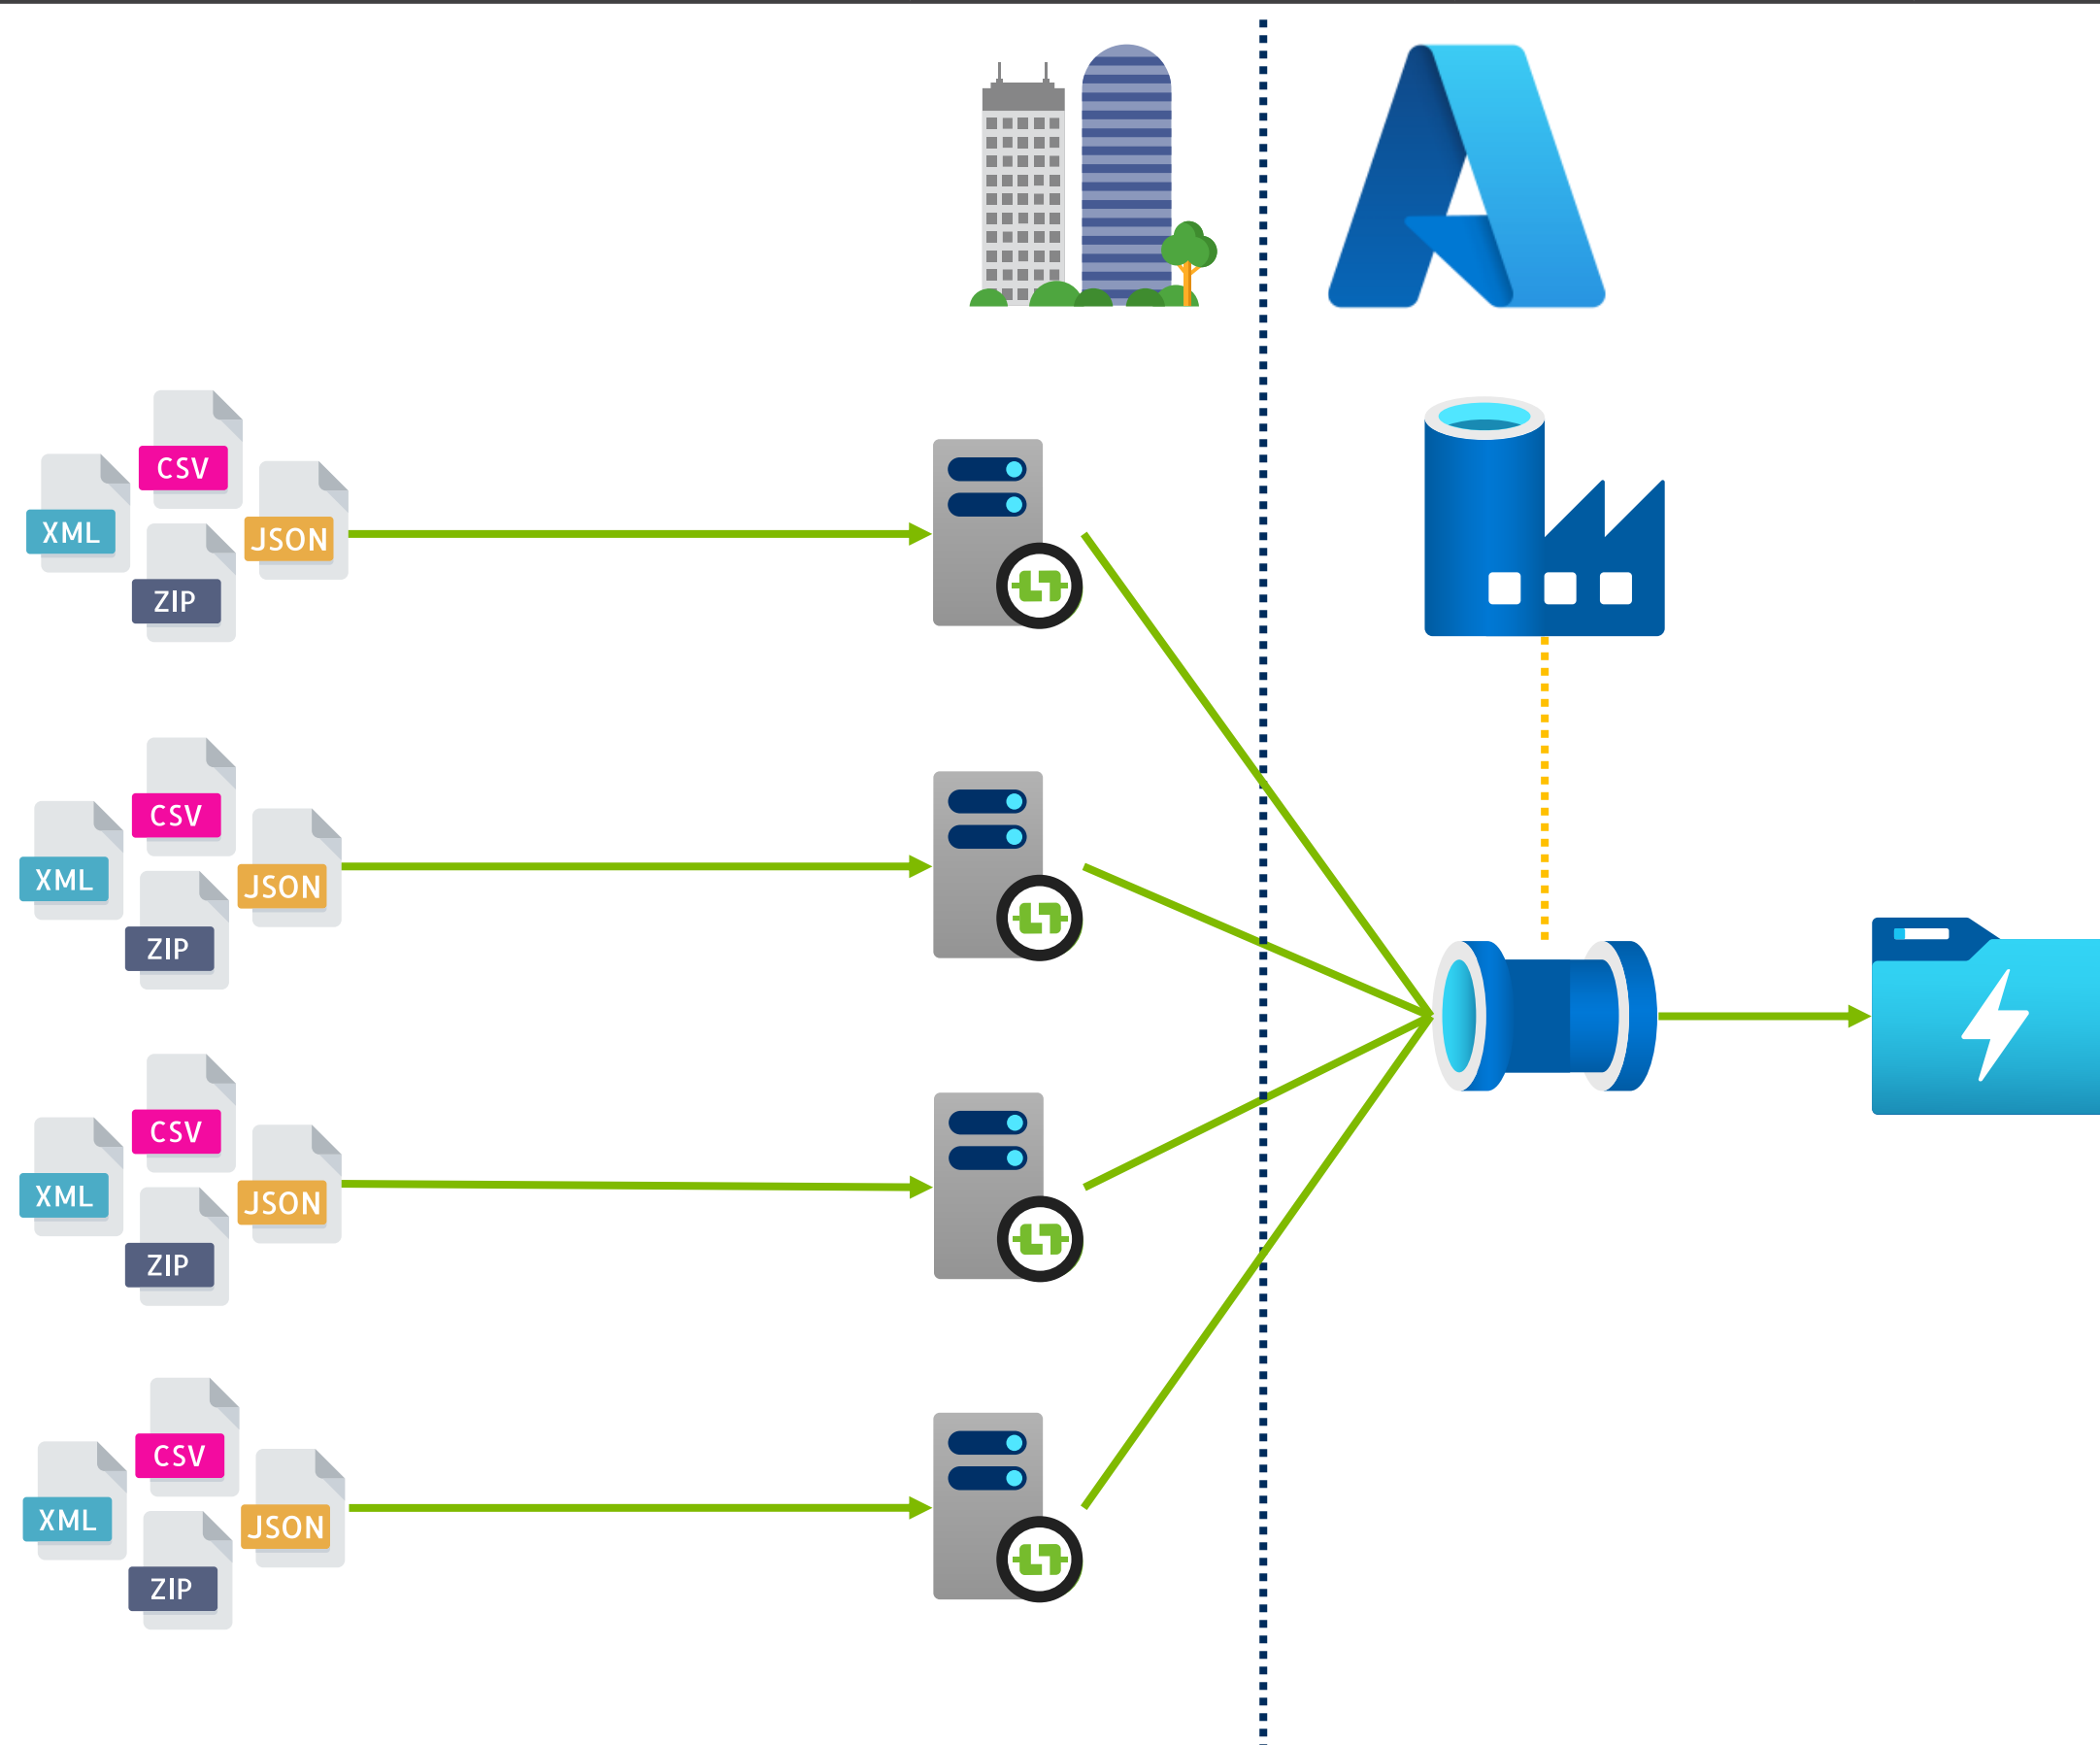



Multiple Hosted IR's at Different Sites








Multiple Hosted IR's as Separate Endpoints










Single Hosted IR Linked to Multiple Data Factories








Hosted IR's with Express Route Connections







Hosted IR Combinations



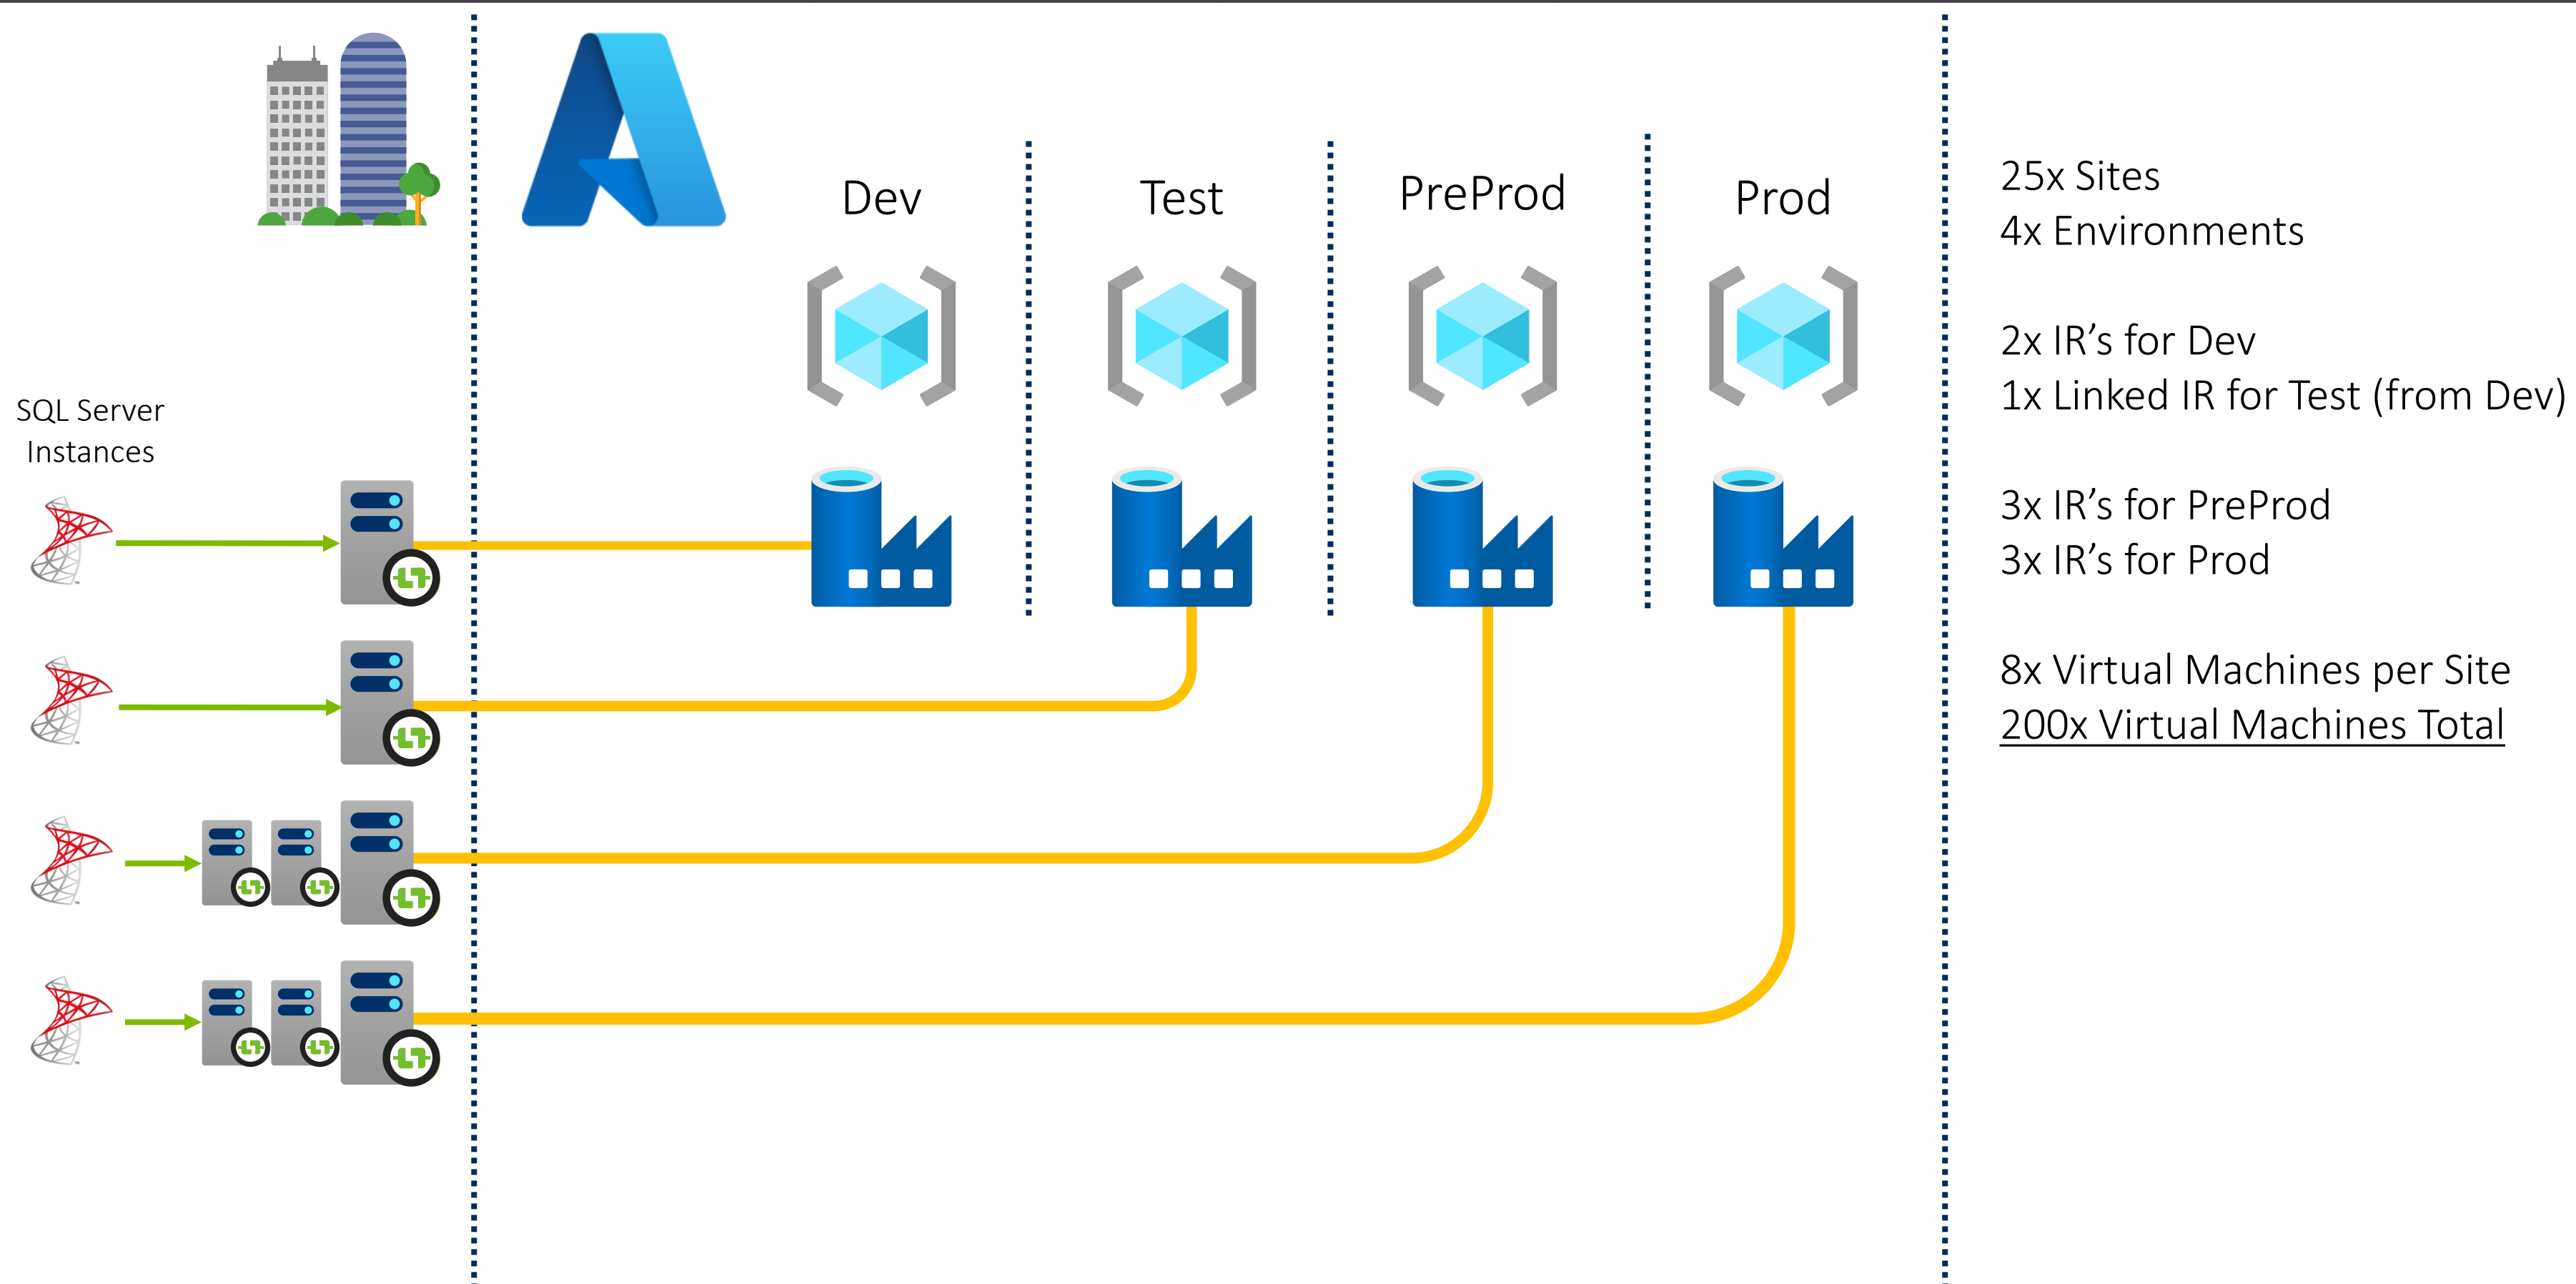


Horizontal vs Vertical Environment Processing



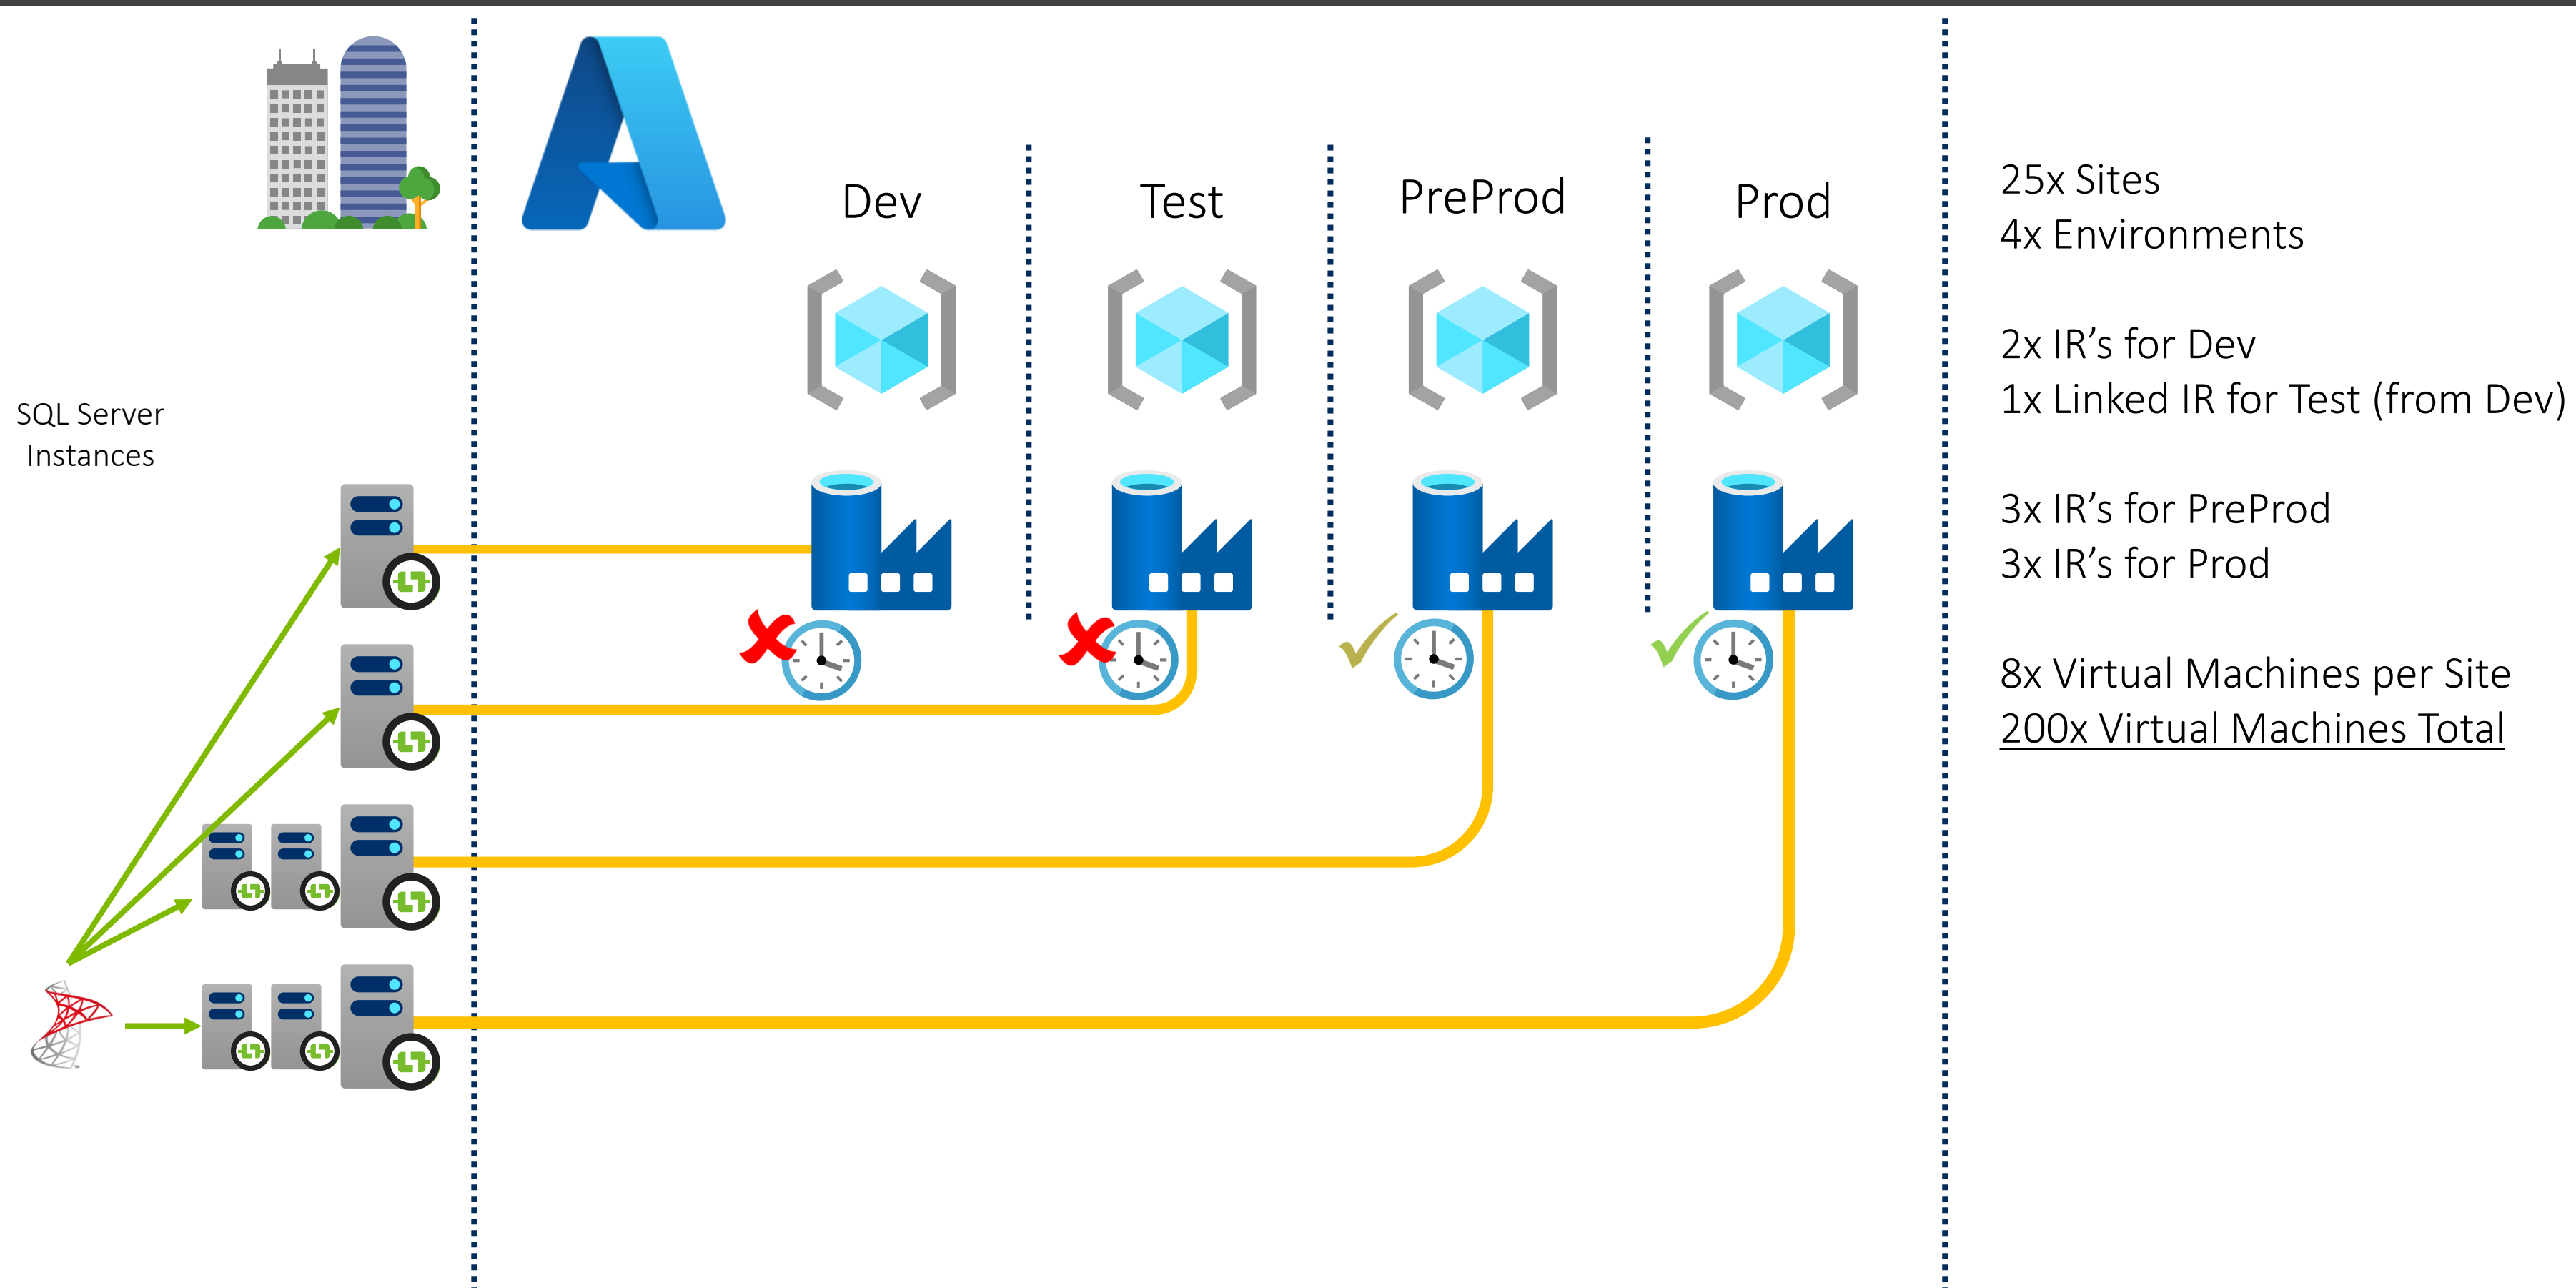


Hosted IR's vs Environments







Hosted IR's vs Environments





Hosted IR's vs Environments



SQL Server
Instances

Dev

Test

PreProd











Prod

25x Sites
4x Environments

2x IR's for Dev
1x Linked IR for Test (from Dev)

3x IR's for PreProd
3x IR's for Prod

8x Virtual Machines per Site
200x Virtual Machines Total

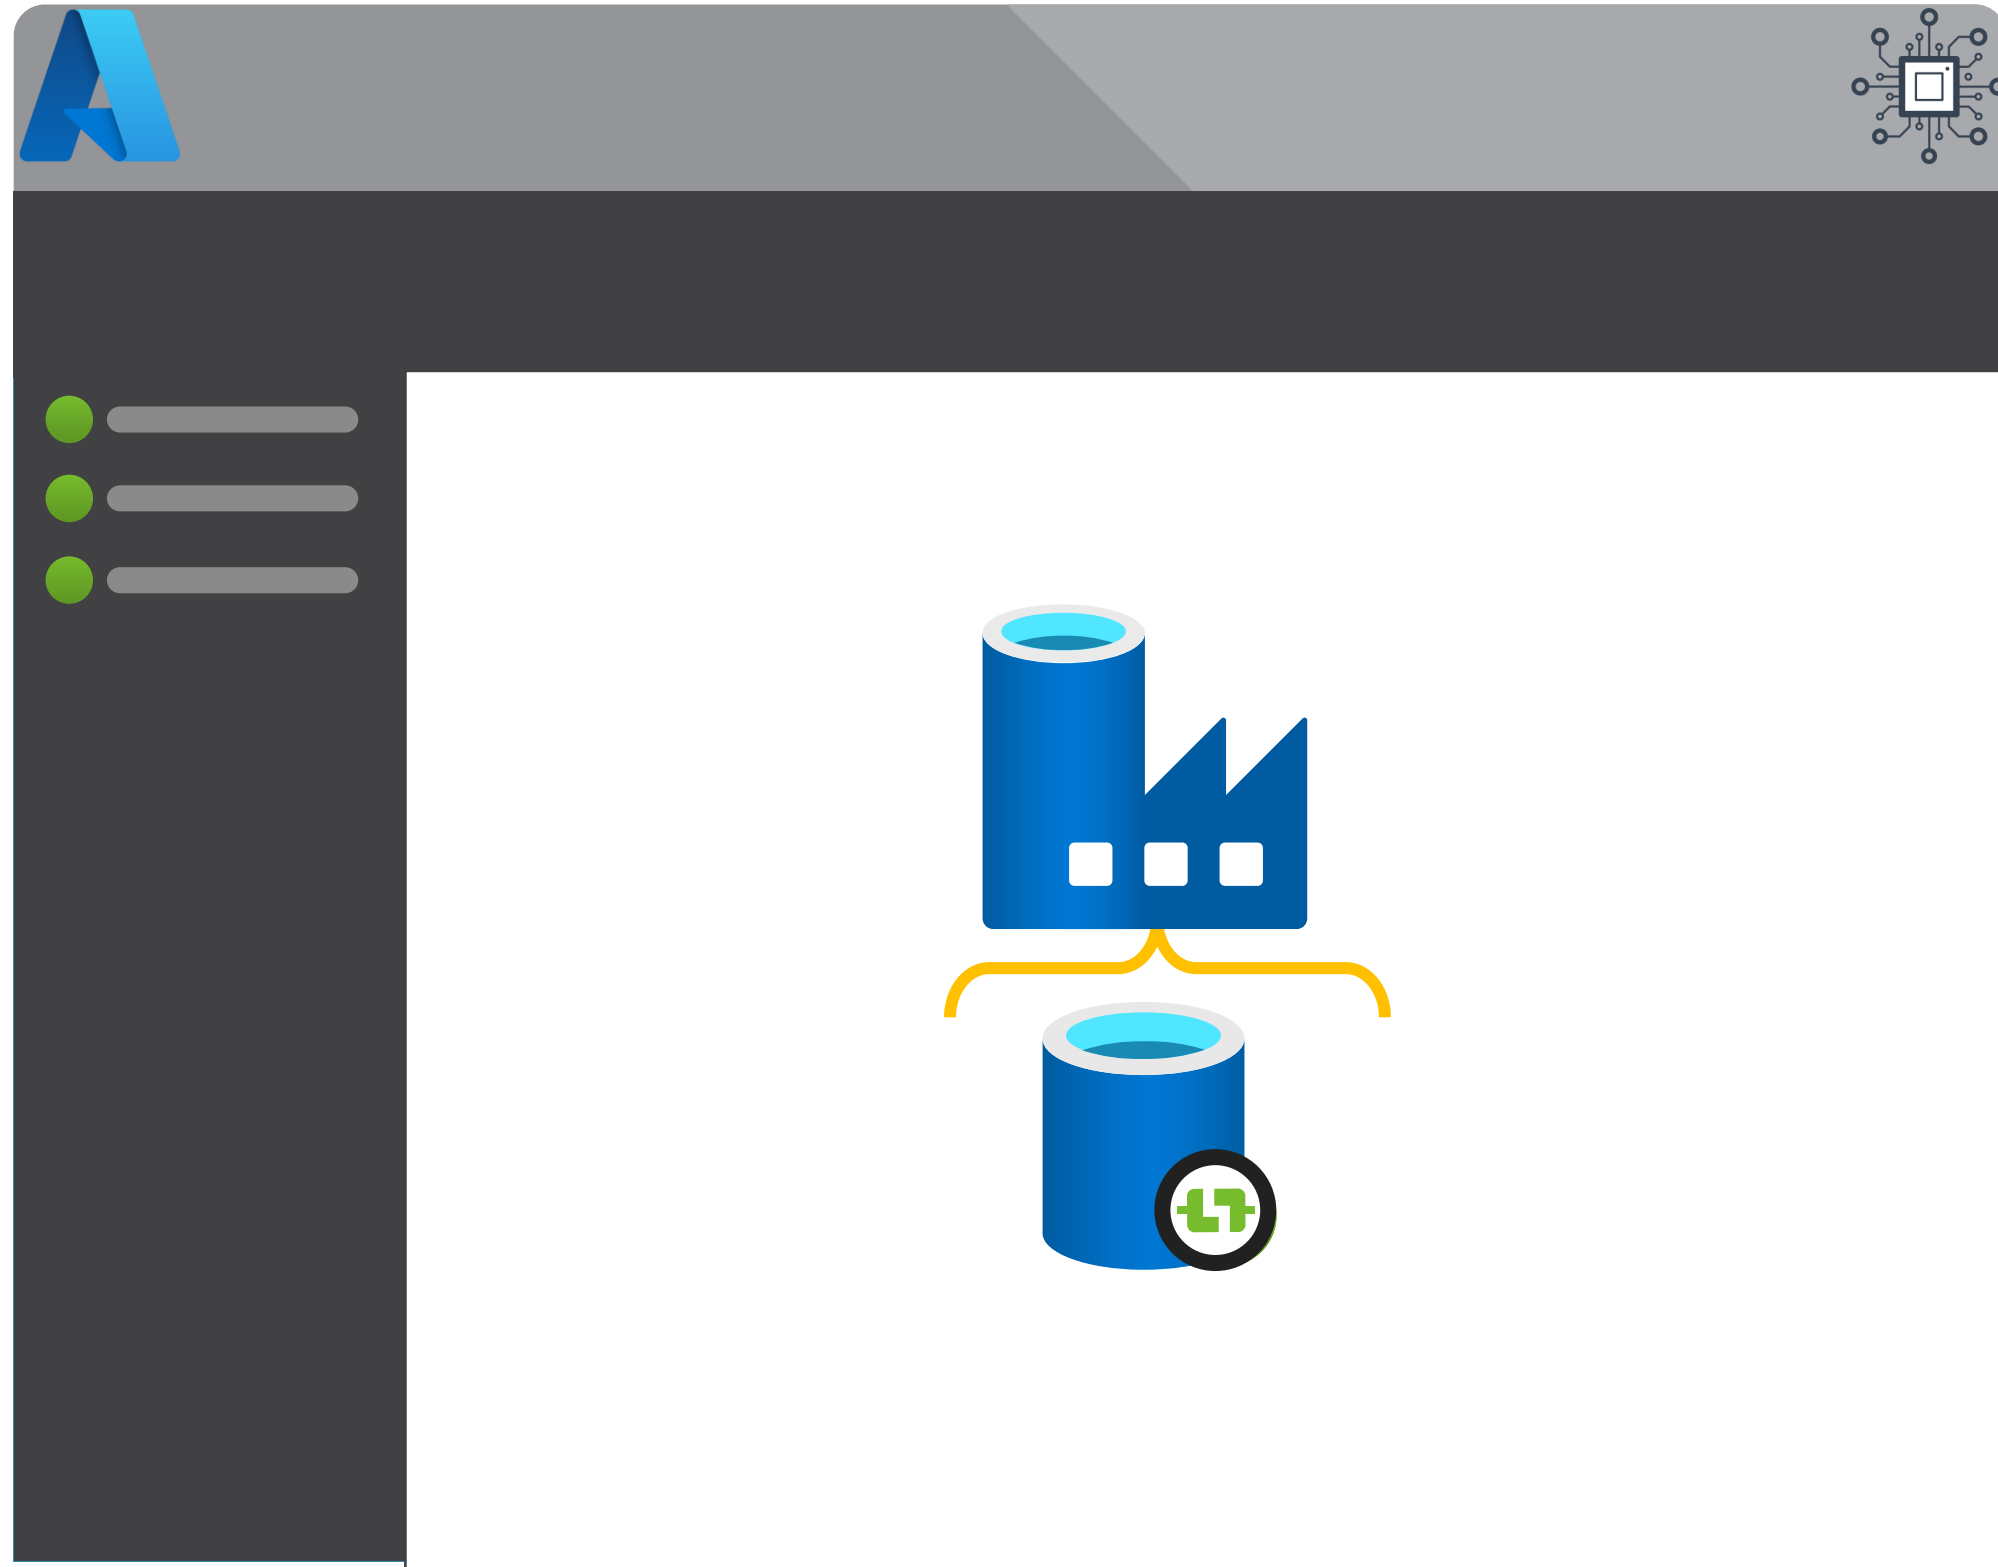


Common



Module 2

Integration Runtime Design Patterns



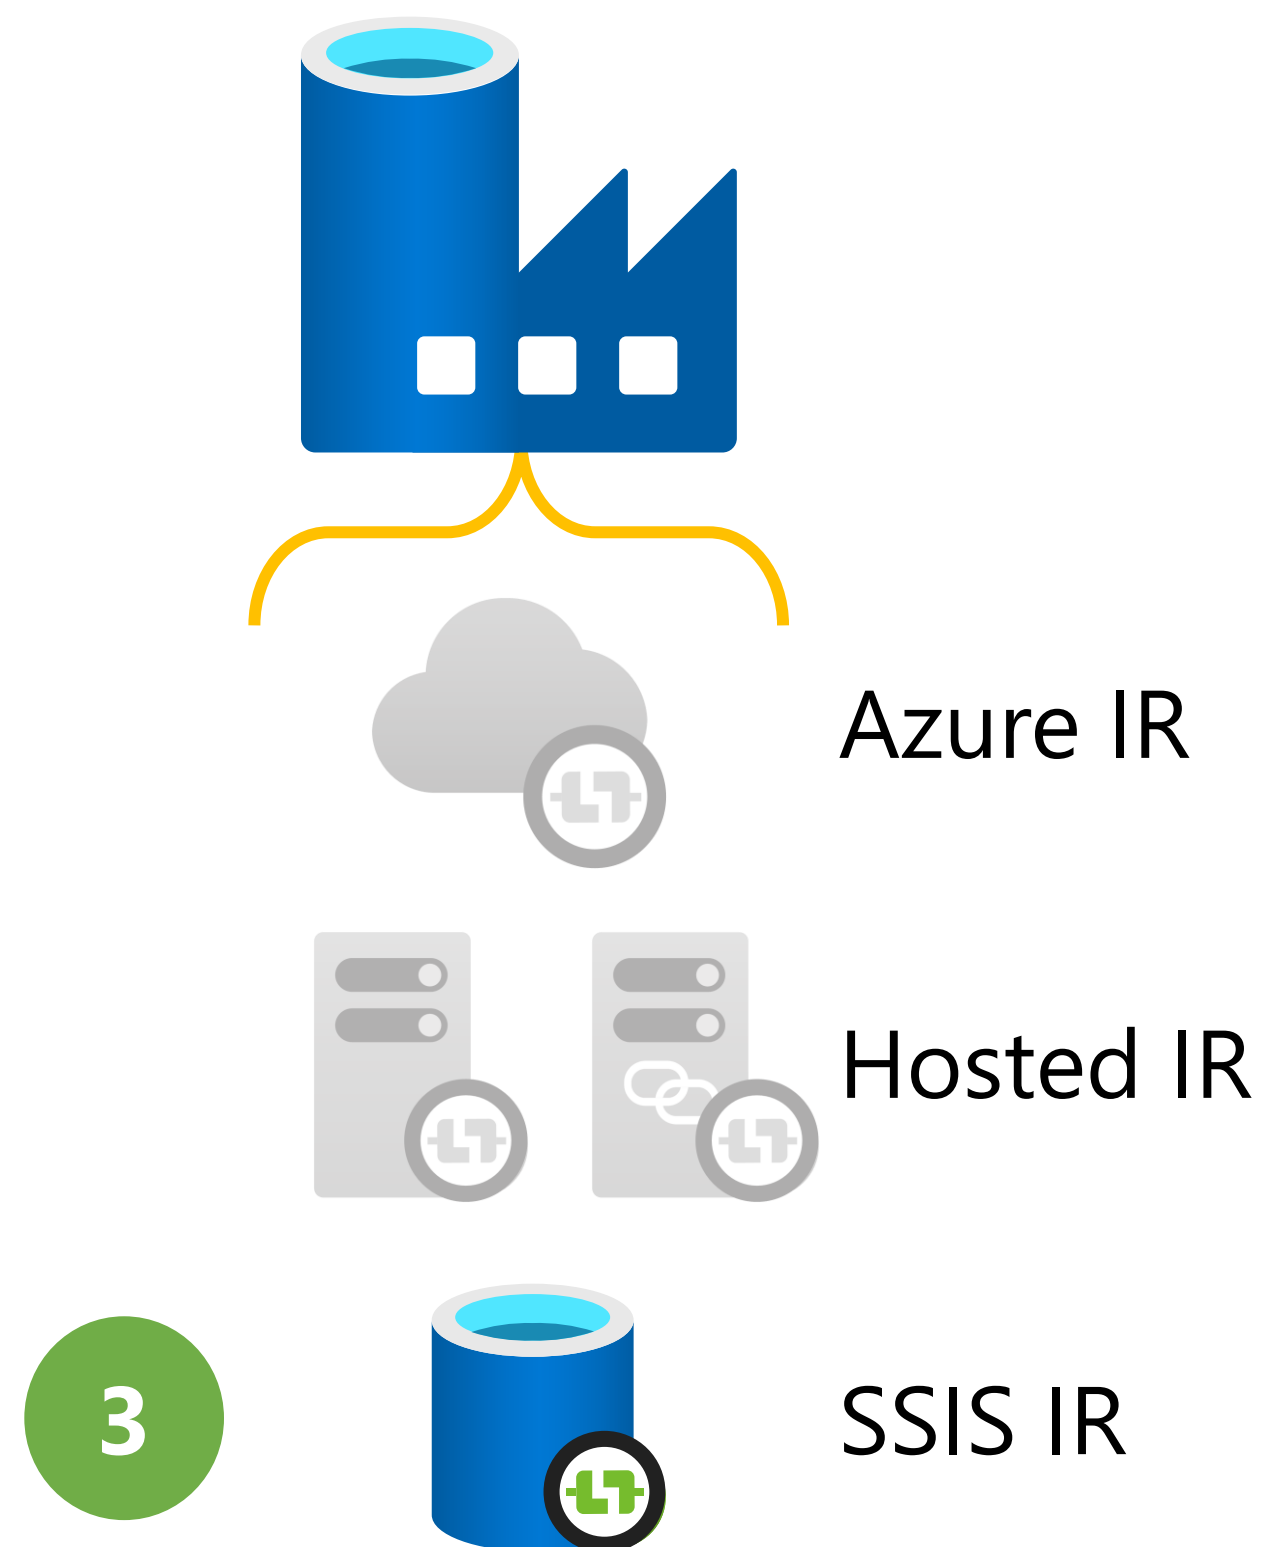

- Compute Types


- Azure
- Hosted
- SSIS

- Patterns & Configuration





What can an Integration Runtime do?





Running an SSIS Package in Azure








SSIS IR



Running an SSIS Package in Azure



SSIS IR







Problem: Using All Of The SSIS IR Compute





SSIS IR



Supports 80 Concurrent Packages
MAXDOP = 80




Runs 1 Package



Parent Package


Child Packages
x80




Pipeline x1




Activities x80



Pipeline x1

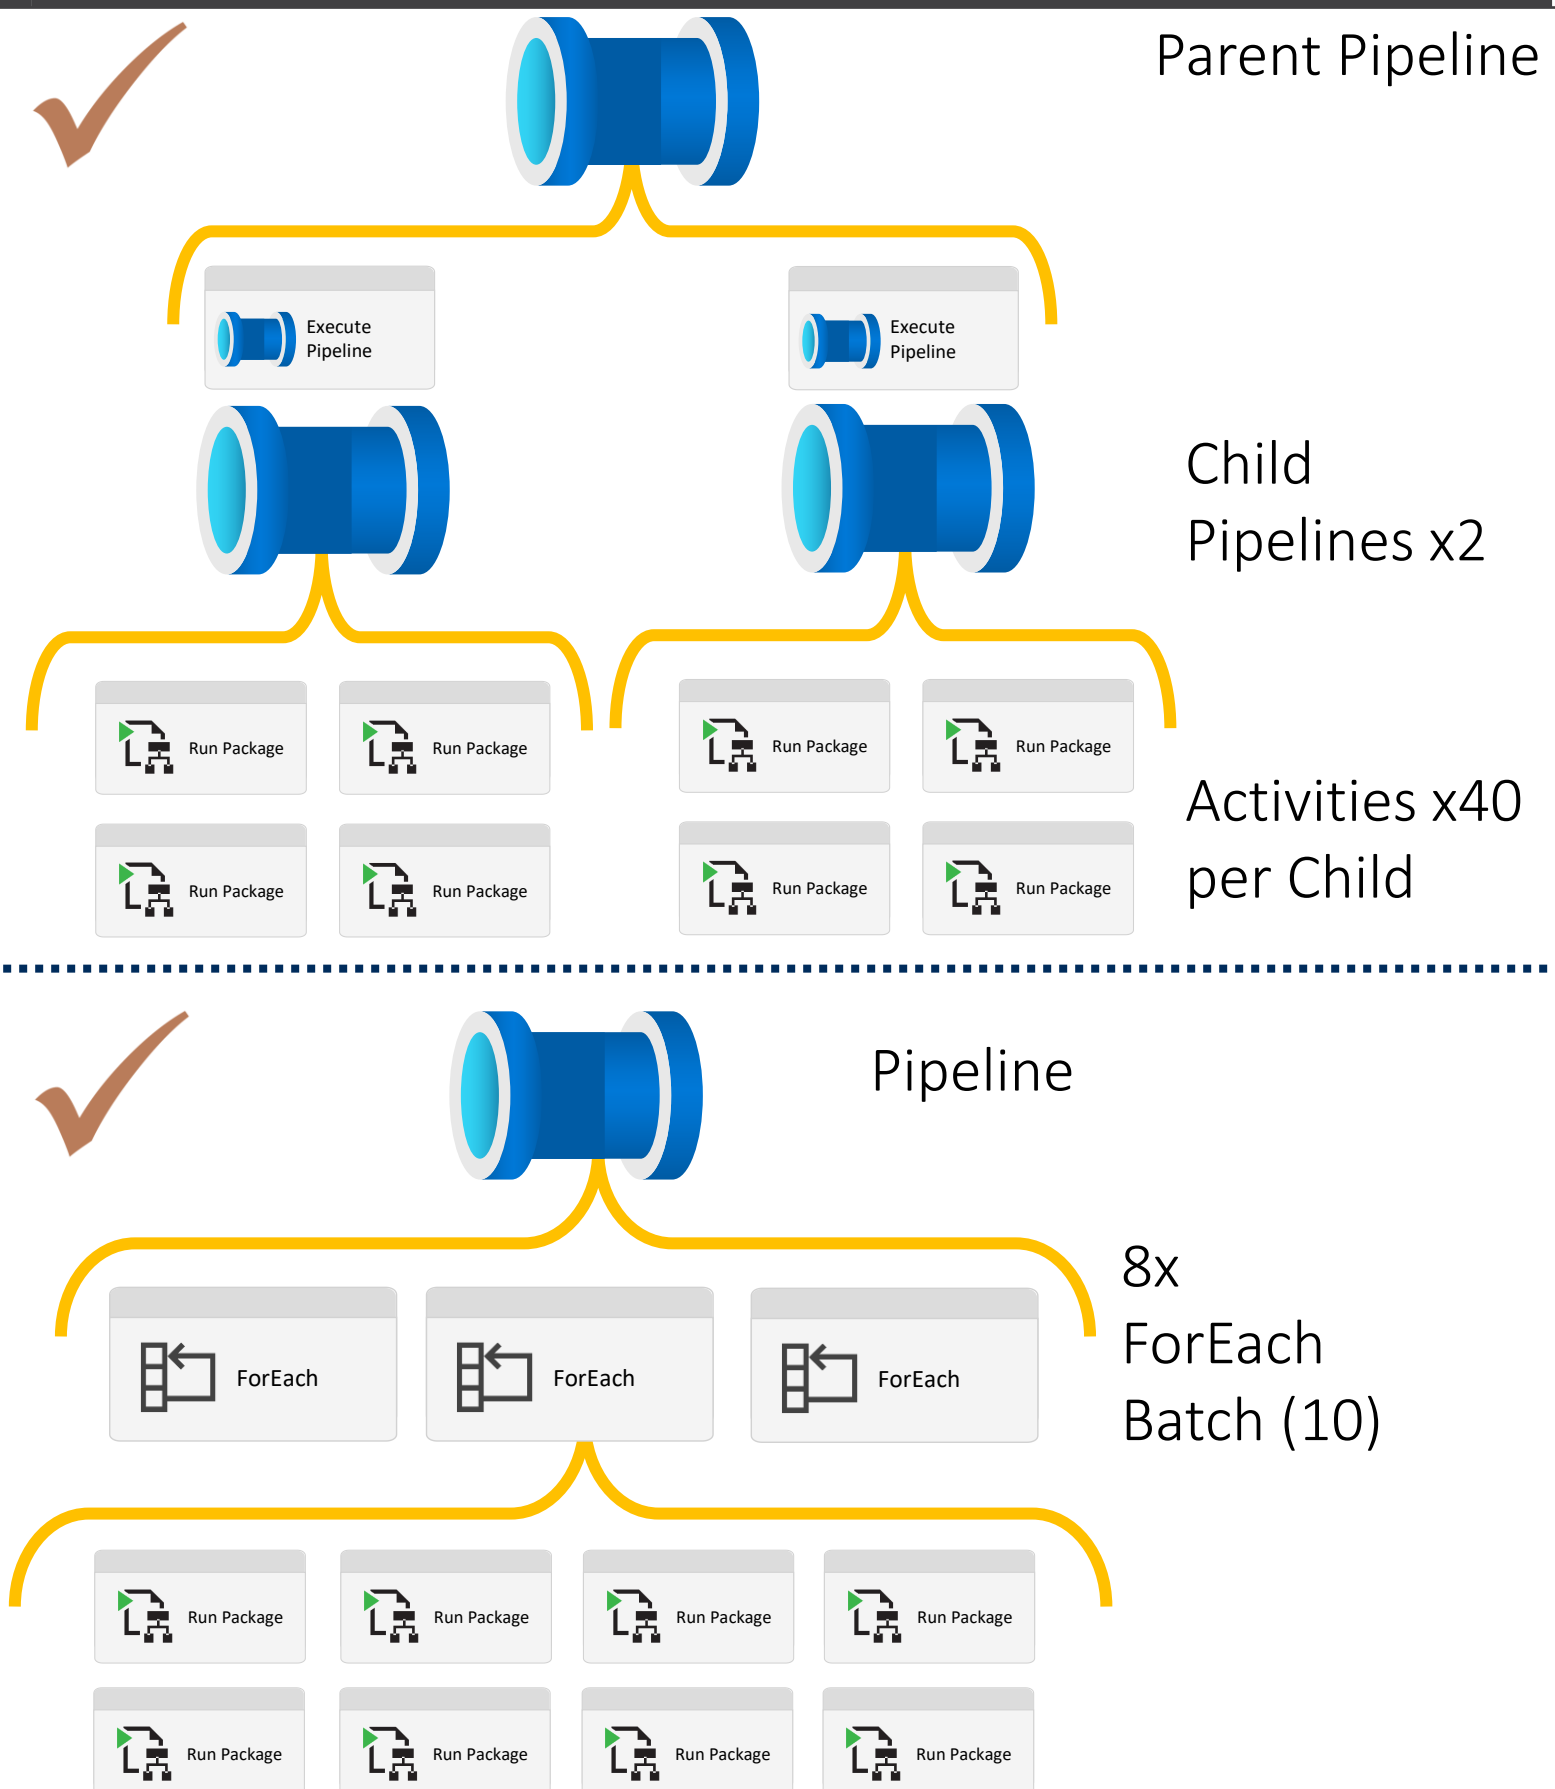
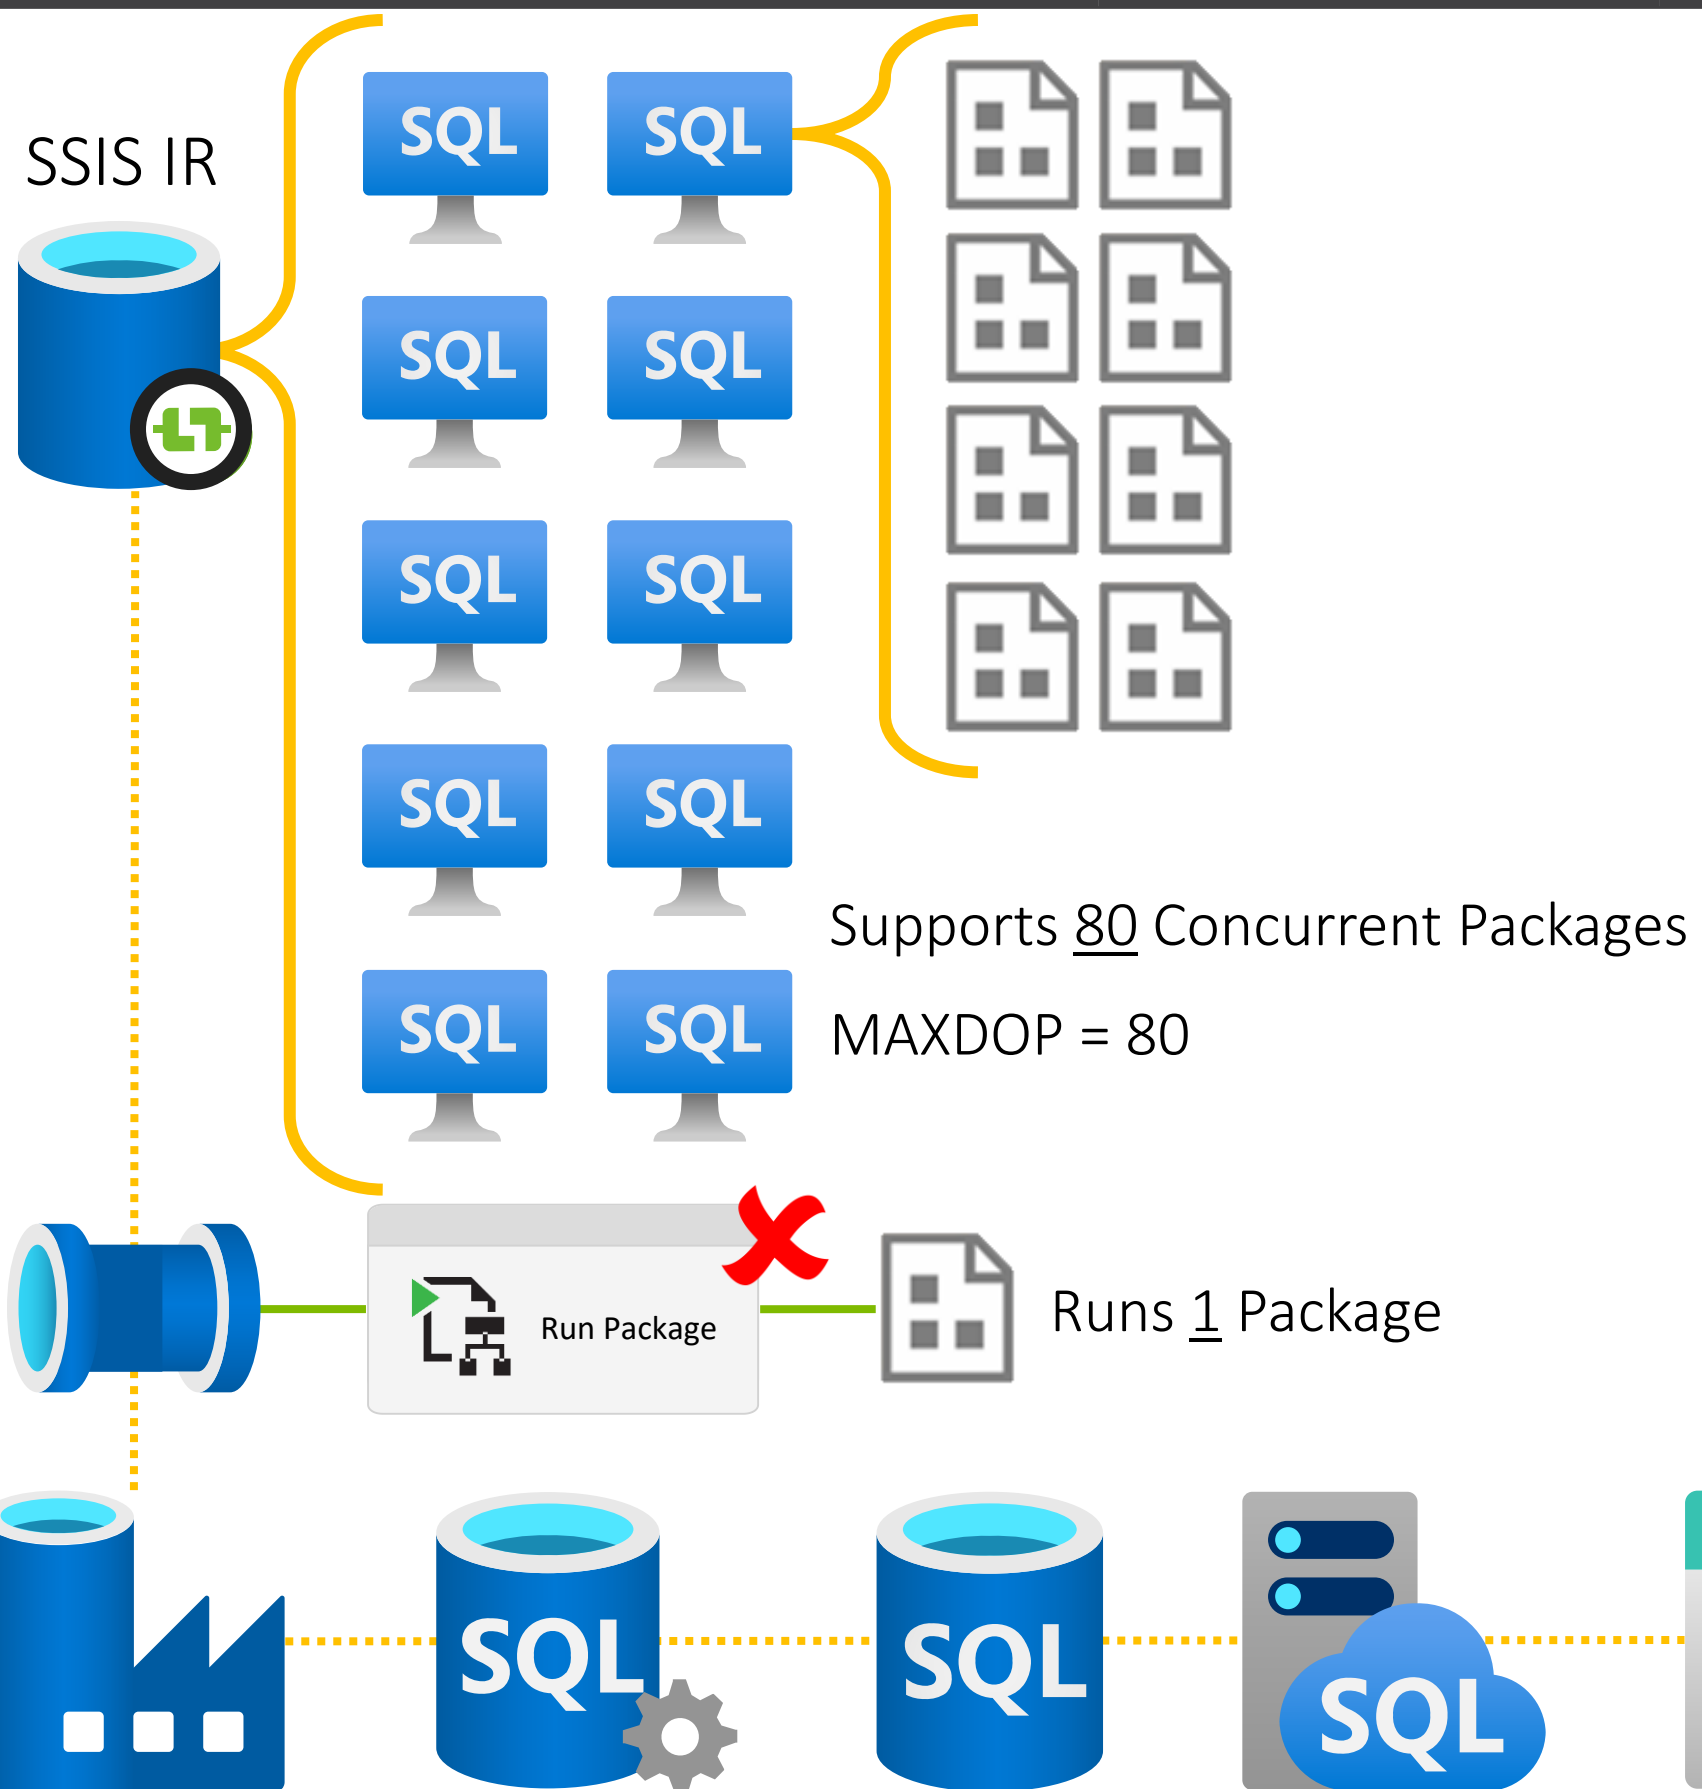


ForEach
Max Batch
(50)





Solution 1 & 2: Static Pipelines











Solution 3: Packages Refactored on a Single Node IR









SSIS IR




Solution 4: Nested ForEach Activities & Bucket Metadata

SSIS IR




FE L1
MAXDOP 50

[Buckets]




@item().BucketId



FE L2
MAXDOP 50




[BucketContents]



Compute size	S4	S6	S7	S9	S12
Max DTUs	200	400	800	1600	3000
Included storage (GB) ¹	250	250	250	250	250
Max storage (GB)	1024	1024	1024	1024	1024
Max in-memory OLTP storage (GB)	N/A	N/A	N/A	N/A	N/A
Max concurrent workers (requests)	400	800	1600	3200	6000
Max concurrent sessions	4800	9600	19200	30000	30000


<https://docs.microsoft.com/en-us/azure/azure-sql/database/resource-limits-dtu-single-databases>

SSISDB




$$(FE L1) \times (FE L2) = NEW MAXDOP$$
$$50 \times 50 = 2500$$




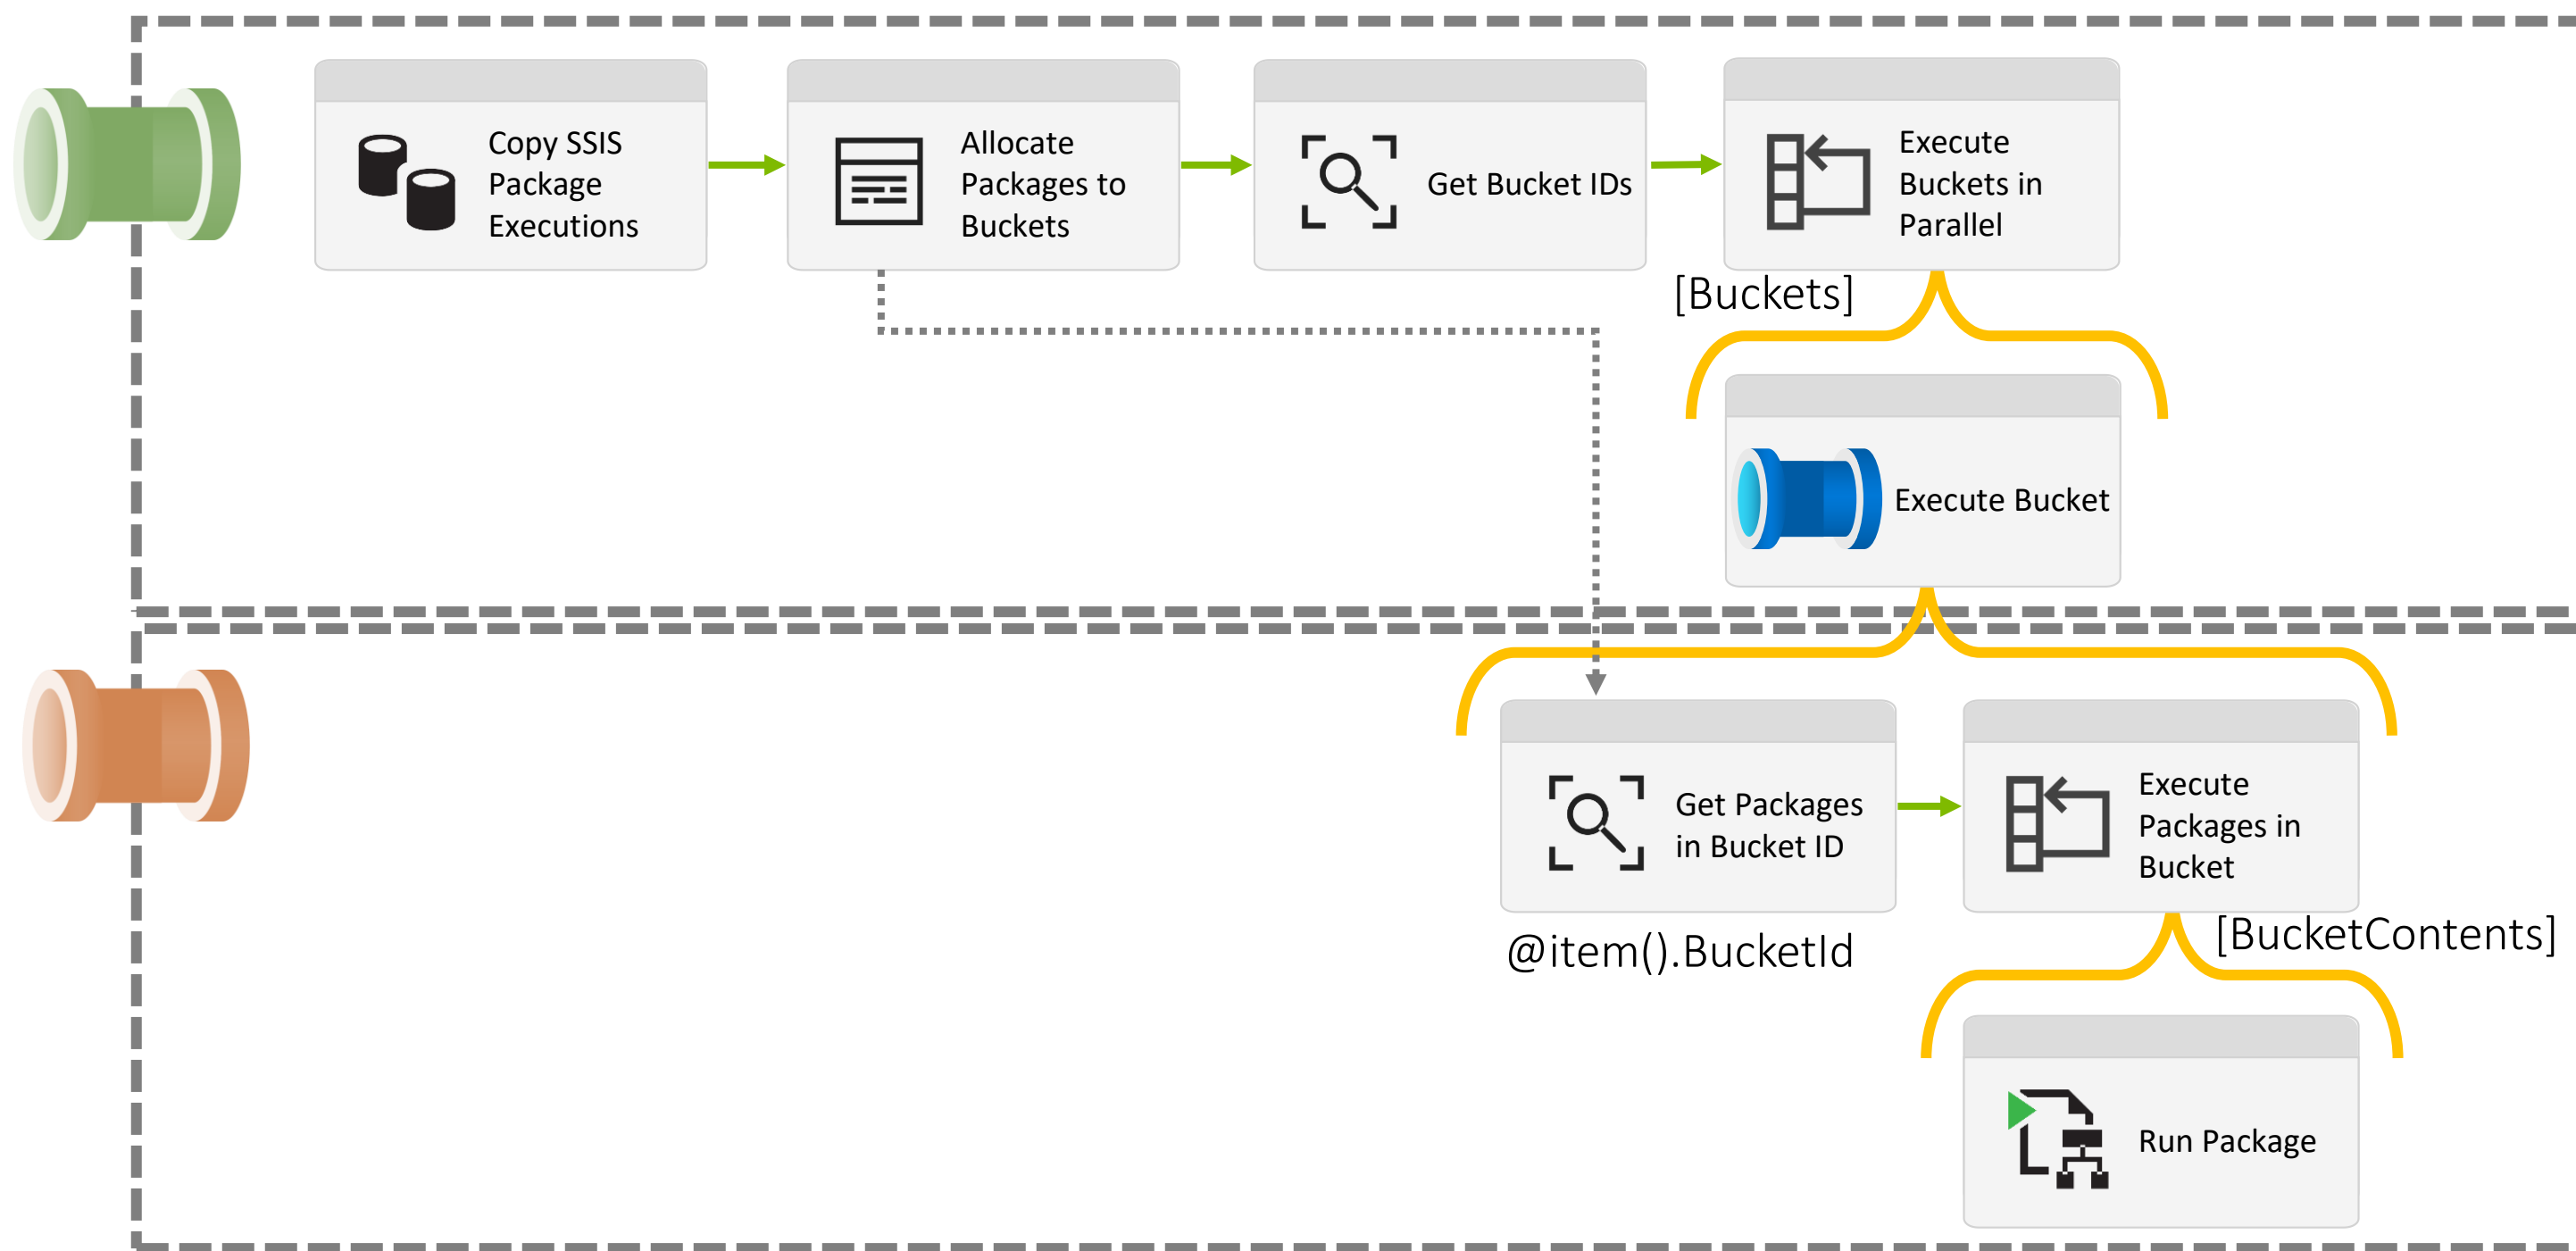





Solution 4: Nested ForEach Activities & Bucket Metadata

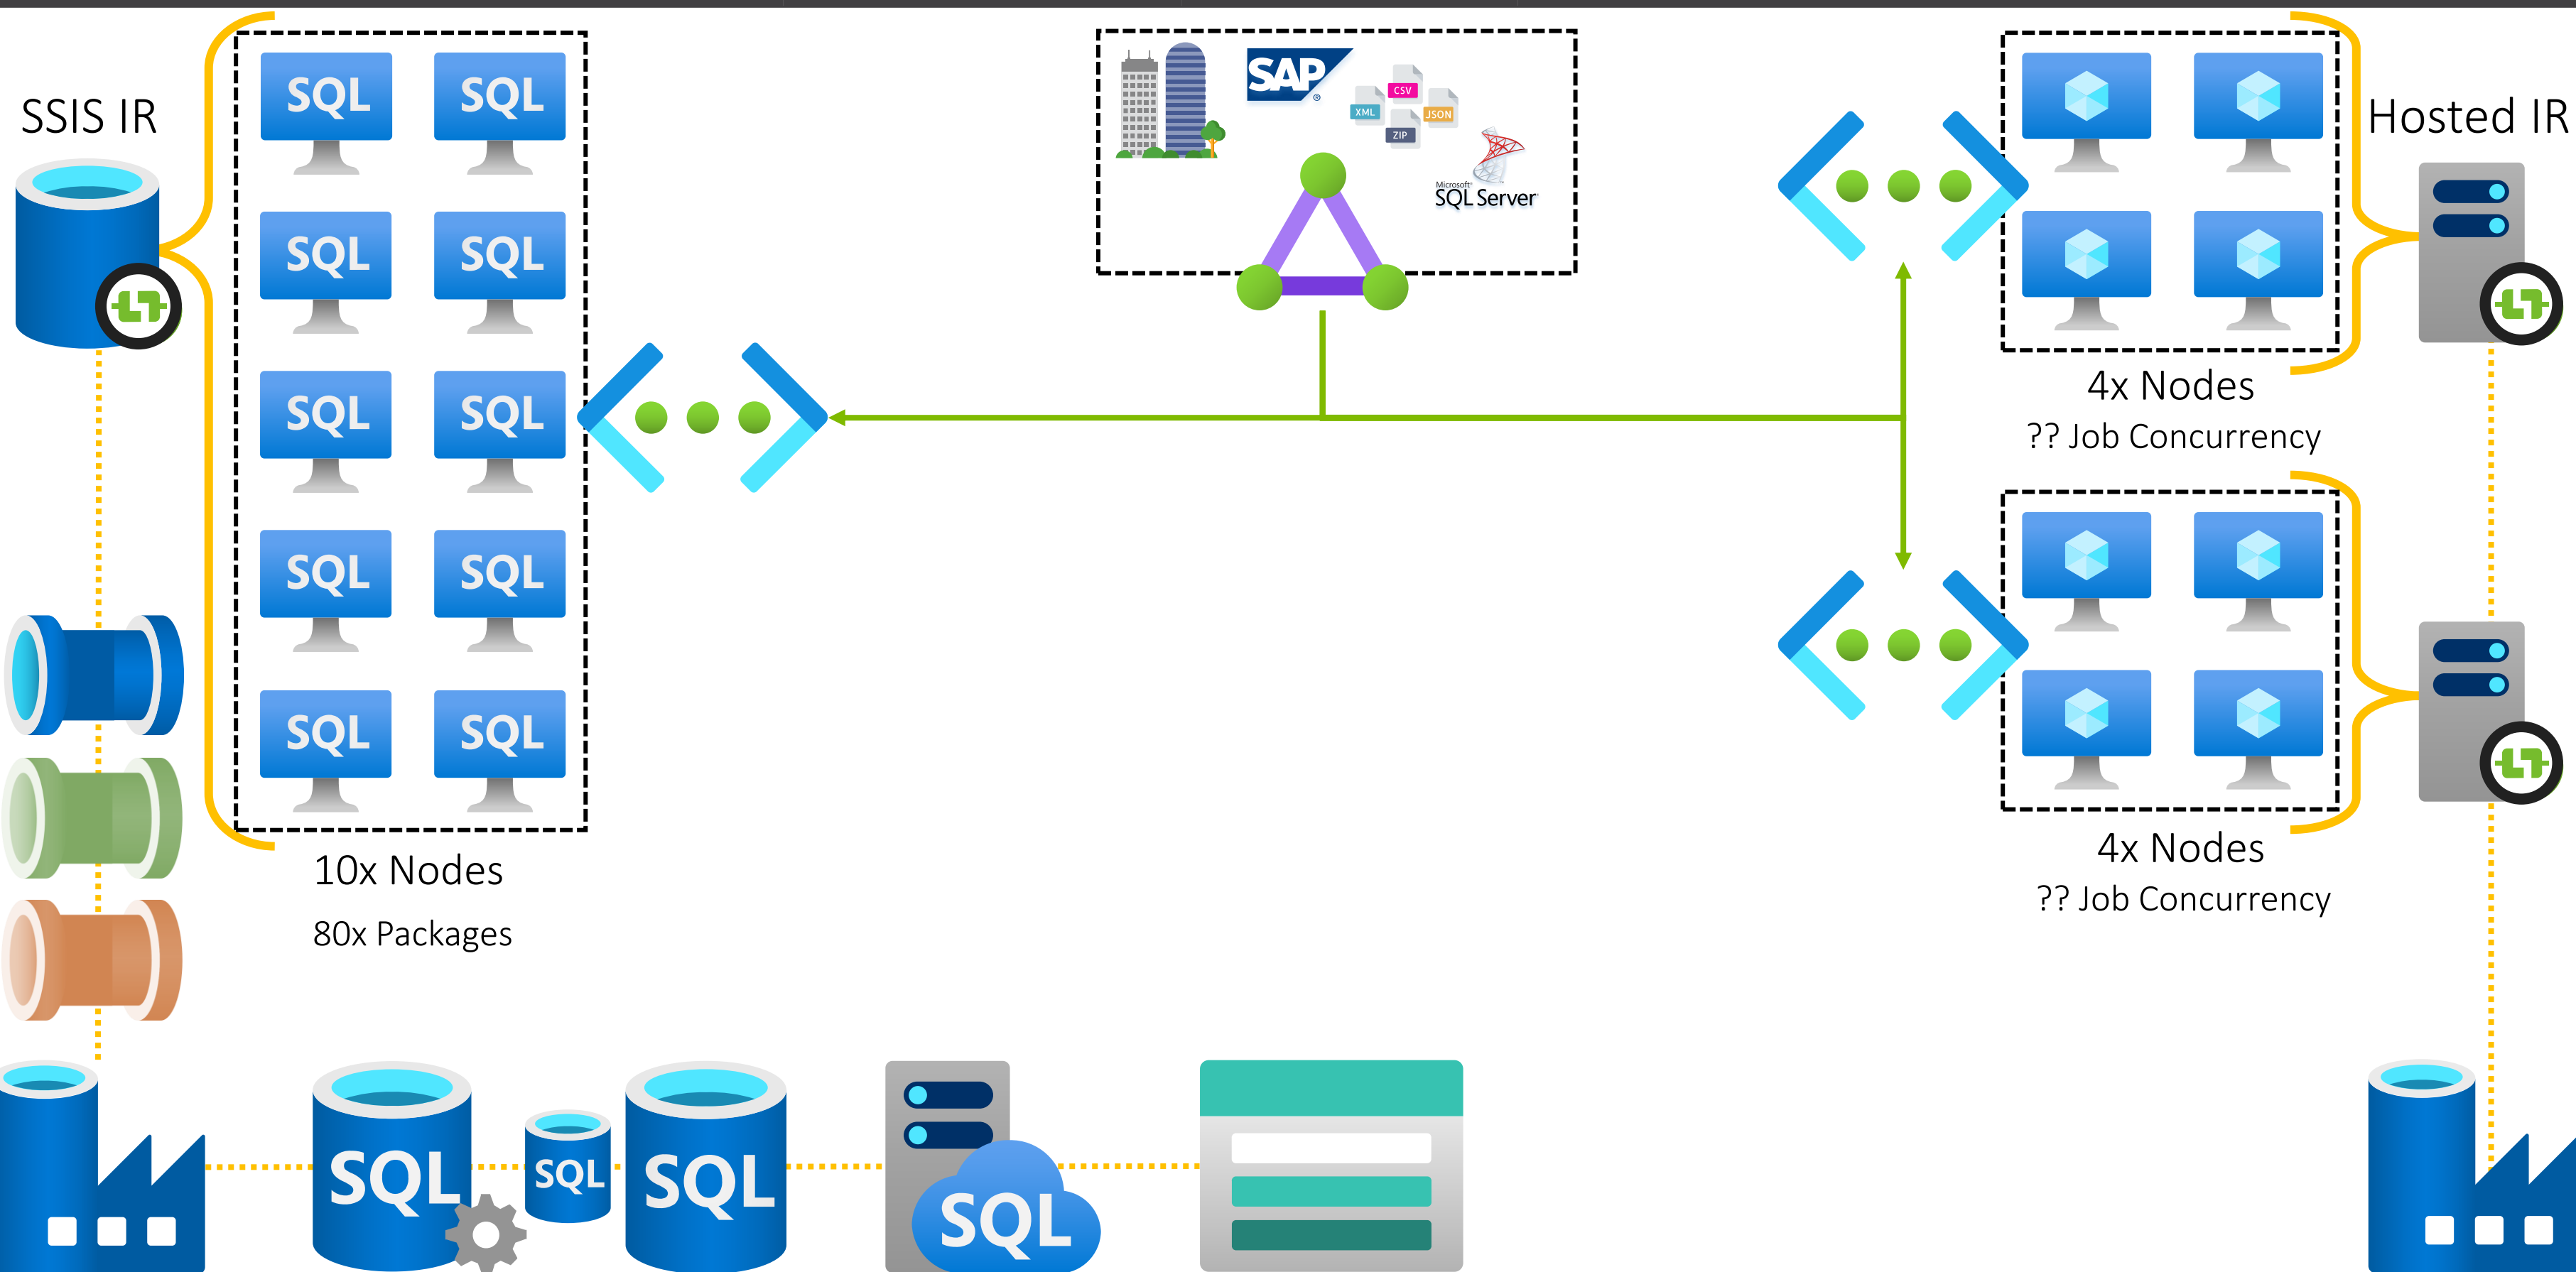



SSIS IR



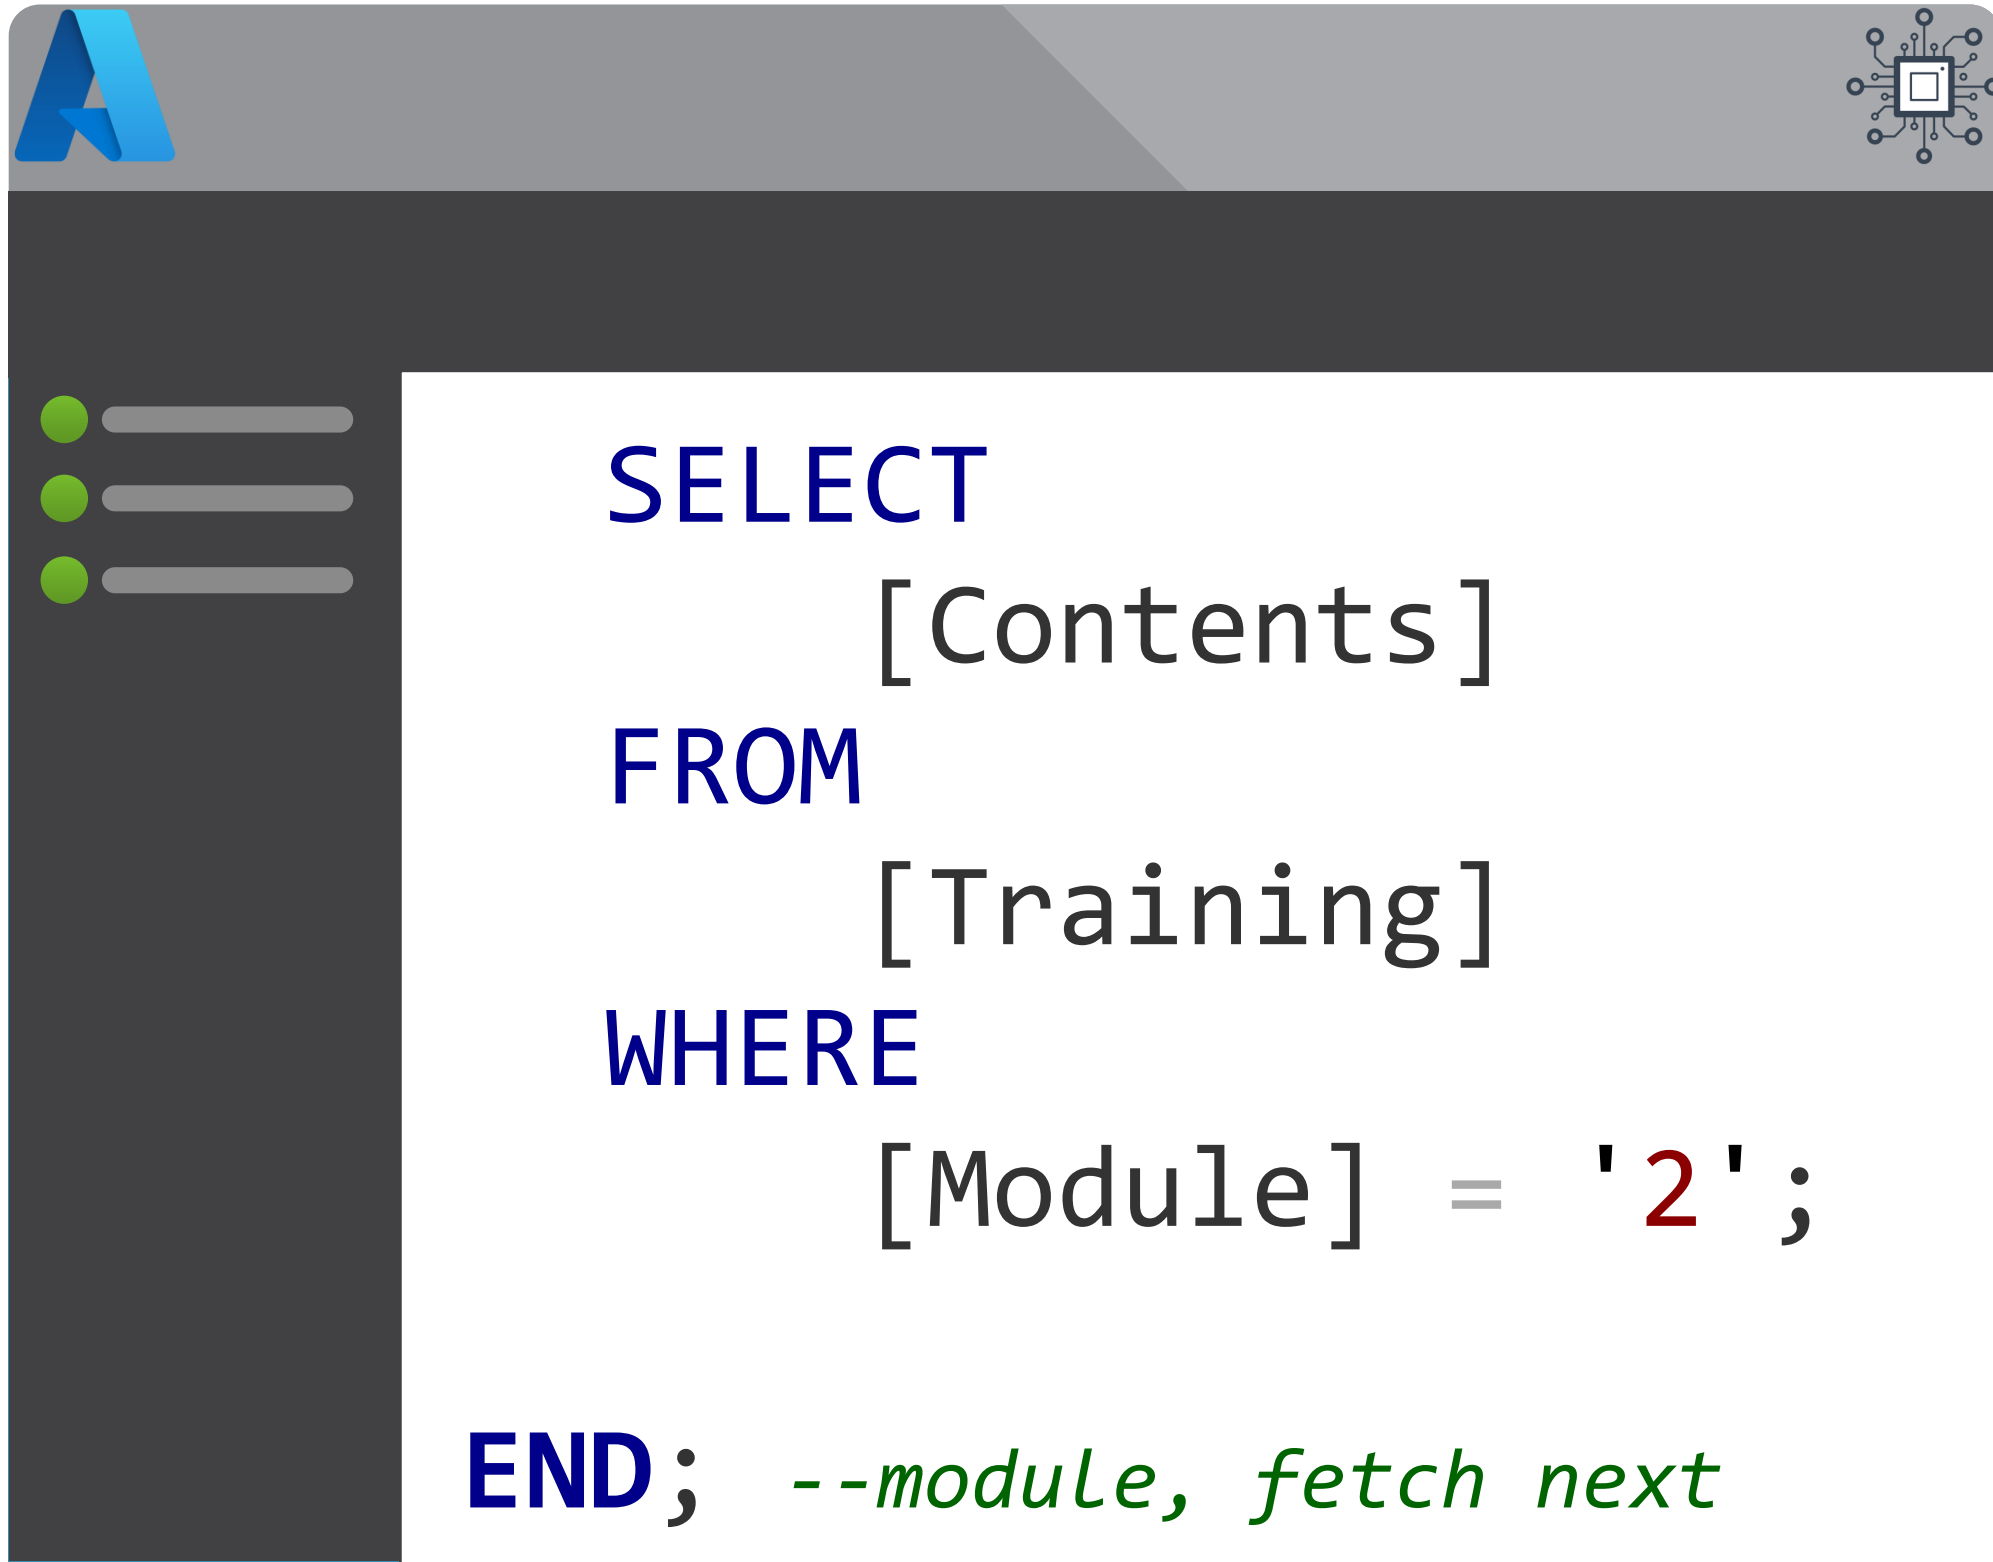


What are the advantages of using this pattern?



Module 2

Integration Runtime Design Patterns



- Compute Types
 - Azure
 - Hosted
 - SSIS
 - AirFlow
- Patterns & Configuration