



Real-Time GIS: Best Practices

Suzanne Foss

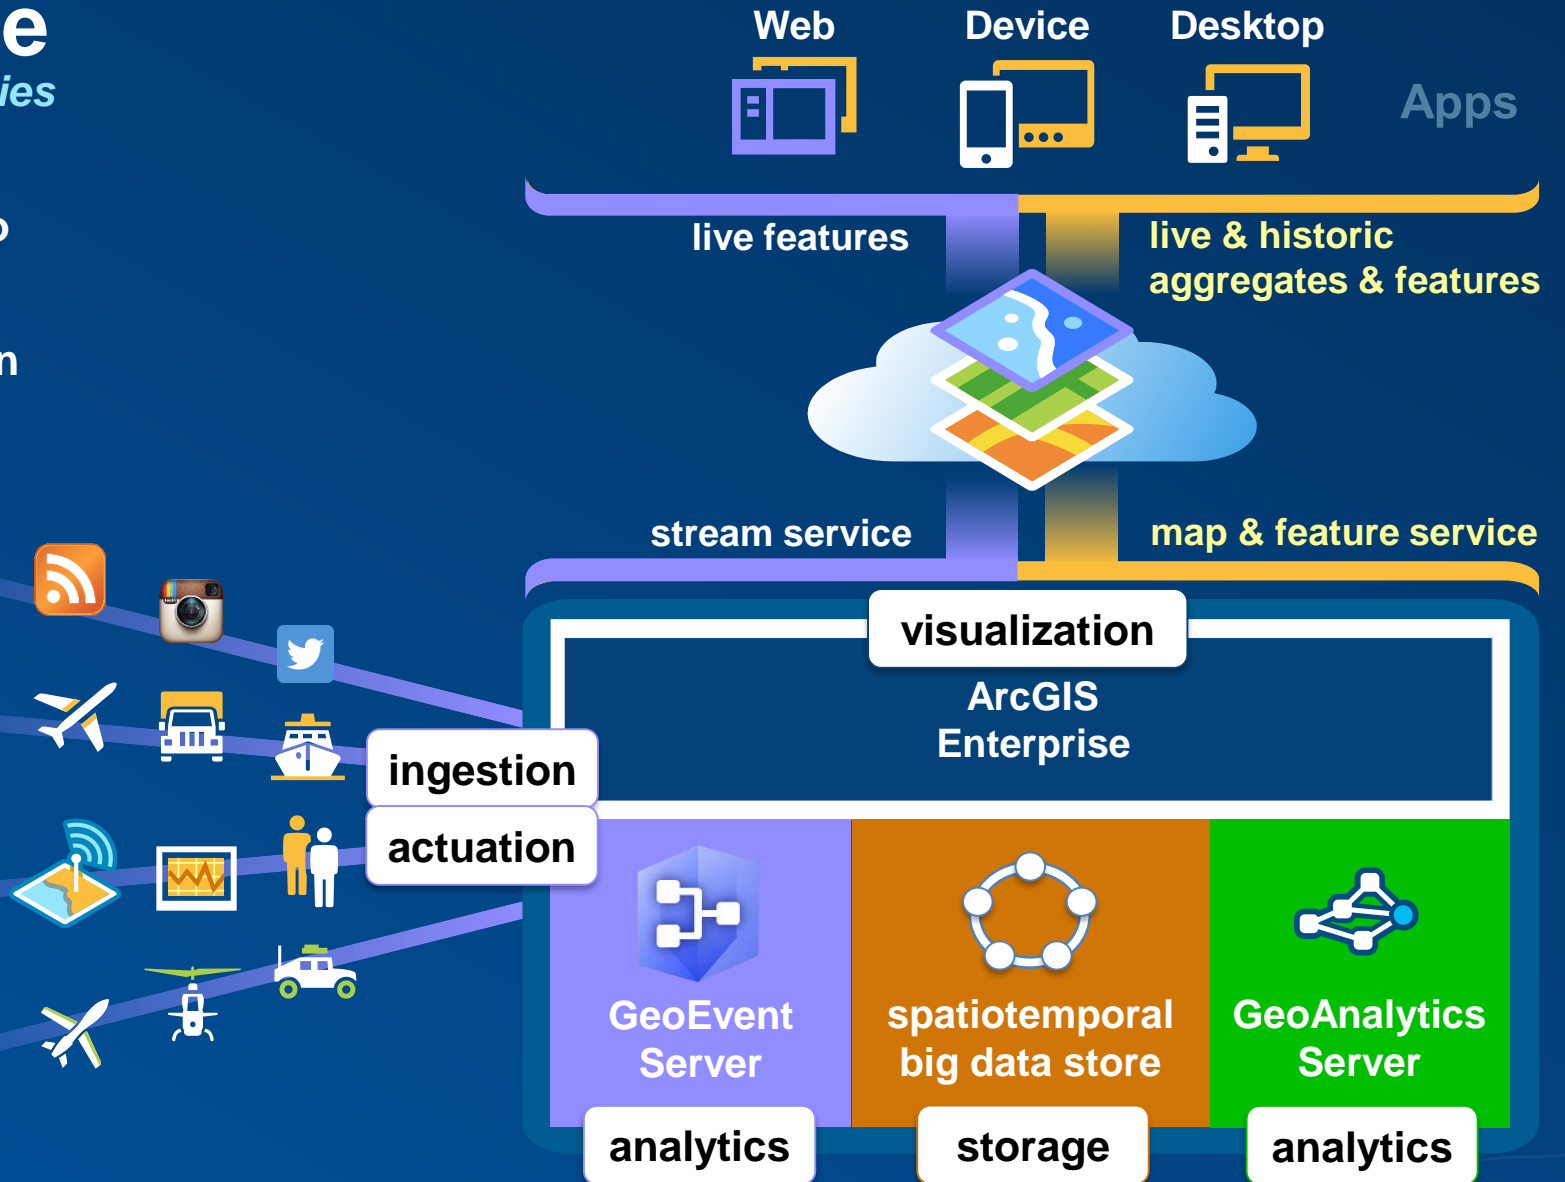
Josh Joyner

2018 Esri Federal GIS Conference | Washington, DC

ArcGIS Enterprise

with real-time & big data capabilities

- Ingest high velocity real-time & Internet of Things (IoT) data into ArcGIS.
- Perform continuous analytics on events as they are received.
- Store observations in a spatiotemporal big data store.
- Visualize high velocity & volume data:
 - as an aggregation
 - or as discrete features.
- Notify about patterns of interest.
- Adjust behavior of things in our environment through actuation.



Agenda:

- 1 Architecture Recommendations
 - 2 Big Data Storage
 - 3 Performance, Resiliency, & Scalability
 - 4 Stream Services
 - 5 Service Design Considerations
 - 6 Upgrade Planning
 - 7 Troubleshooting
-



Architecture Recommendations

GeoEvent Server

What are the primary factors I should consider?

- **Operating environment:** *m4.2xlarge*
 - virtual machines – beware! resources need to be shared in an effective way, like EC2 or Azure.
 - dedicated bare metal machines or public cloud instances are much more deterministic.
- **Network** *1 GBit/s*
 - speed – the faster the better.
- **Memory** *32GiB, default JVM max heap size is 4 GB*
 - size – 8GB has been required since 10.3.
 - type – minimum of DDR3 is recommended.
 - clock speed (MHz) and transfer rate (Mbps) – the faster the better.
- **Processors** *8 vCPU*
 - # of cores – the more the better.
 - speed (GHz) – the faster the better.
- **Disk** *10GB recommended minimum (new for 10.6)*
 - **700MB required for installation**
 - amount of disk space needed will vary based on quantity of deployed input connectors
 - each input can utilize up to a maximum of 600 MB of disk space before clean up

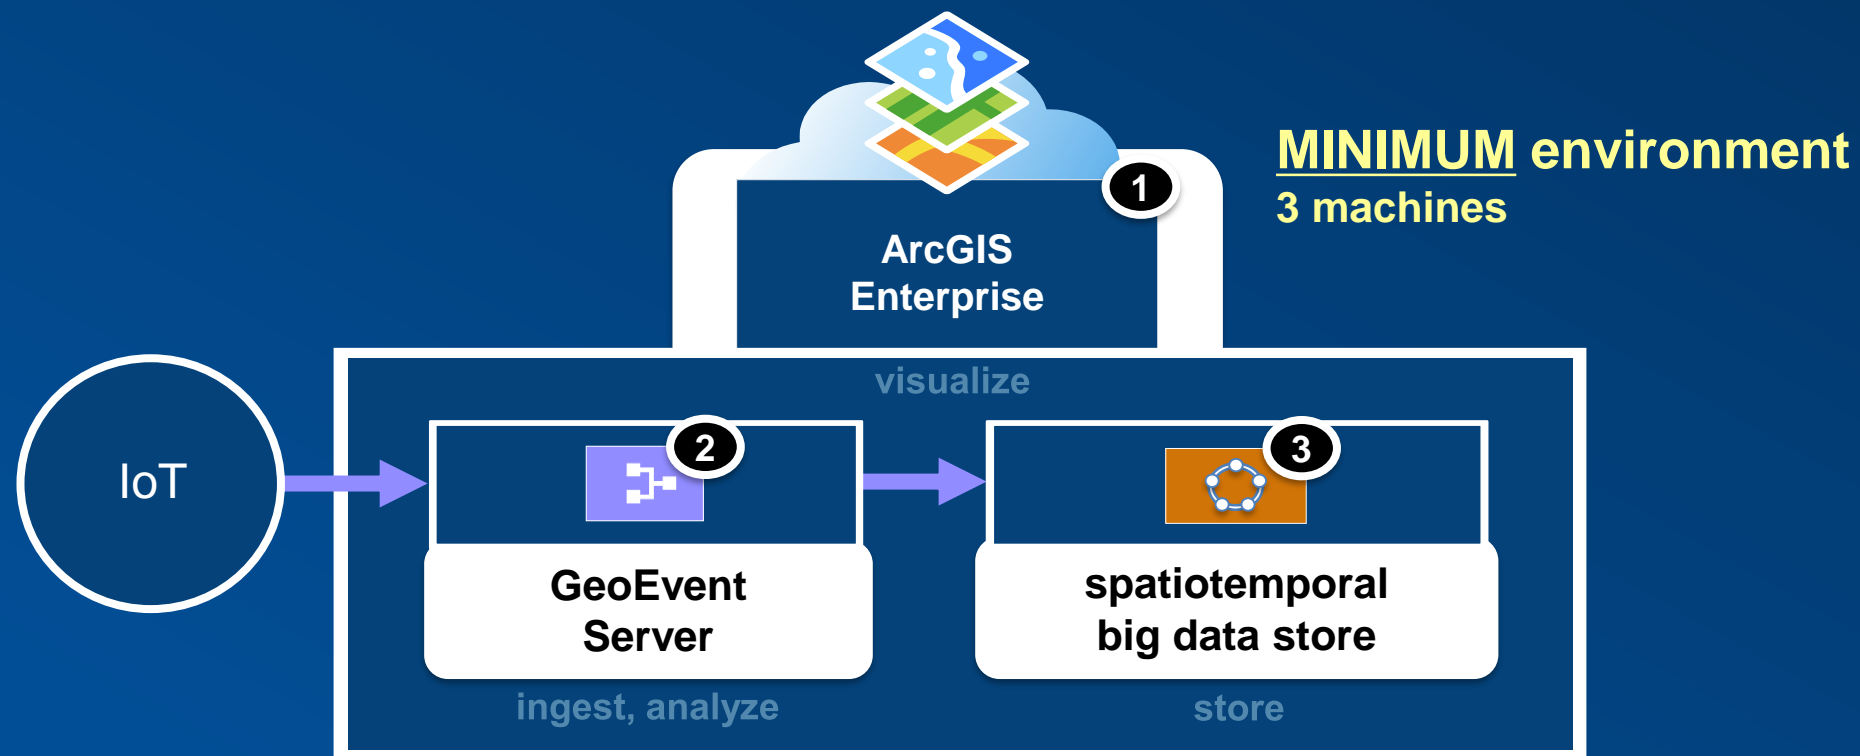
spatiotemporal big data store

What are the primary factors I should consider?

- **Operating environment:** *m4.2xlarge*
 - virtual machines – beware! resources need to be shared in an effective way, like EC2 or Azure.
 - dedicated bare metal machines or public cloud instances are much more deterministic.
- **Disk**
 - speed – the faster the better *1,000 Mbps EBS, note: local SSD is much better*
- **Network**
 - speed – the faster the better. *1 GBit/s*
- **Memory**
 - size – 16GB minimum. *32GiB, big data store allocates 8GiB by default*
 - type – DDR3 is recommended.
 - clock speed (MHz) and transfer rate (Mbps) – the faster the better.
- **Processors**
 - # of cores – the more the better. *8 vCPU*
 - speed (GHz) – the faster the better.

ArcGIS Enterprise

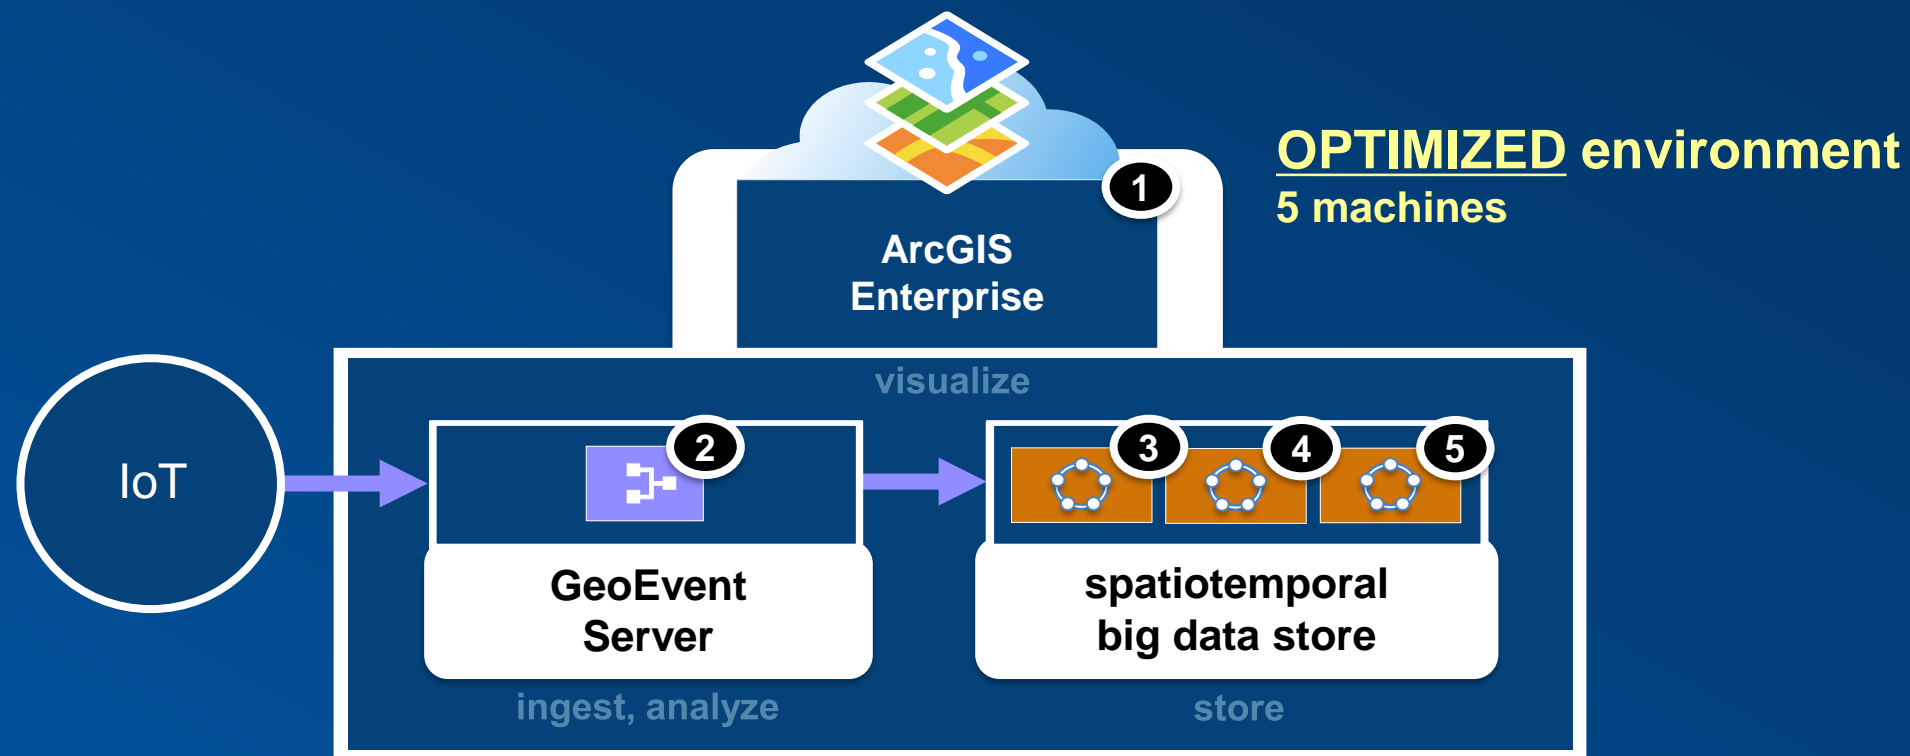
with real-time capabilities



MINIMUM environment
3 machines

ArcGIS Enterprise

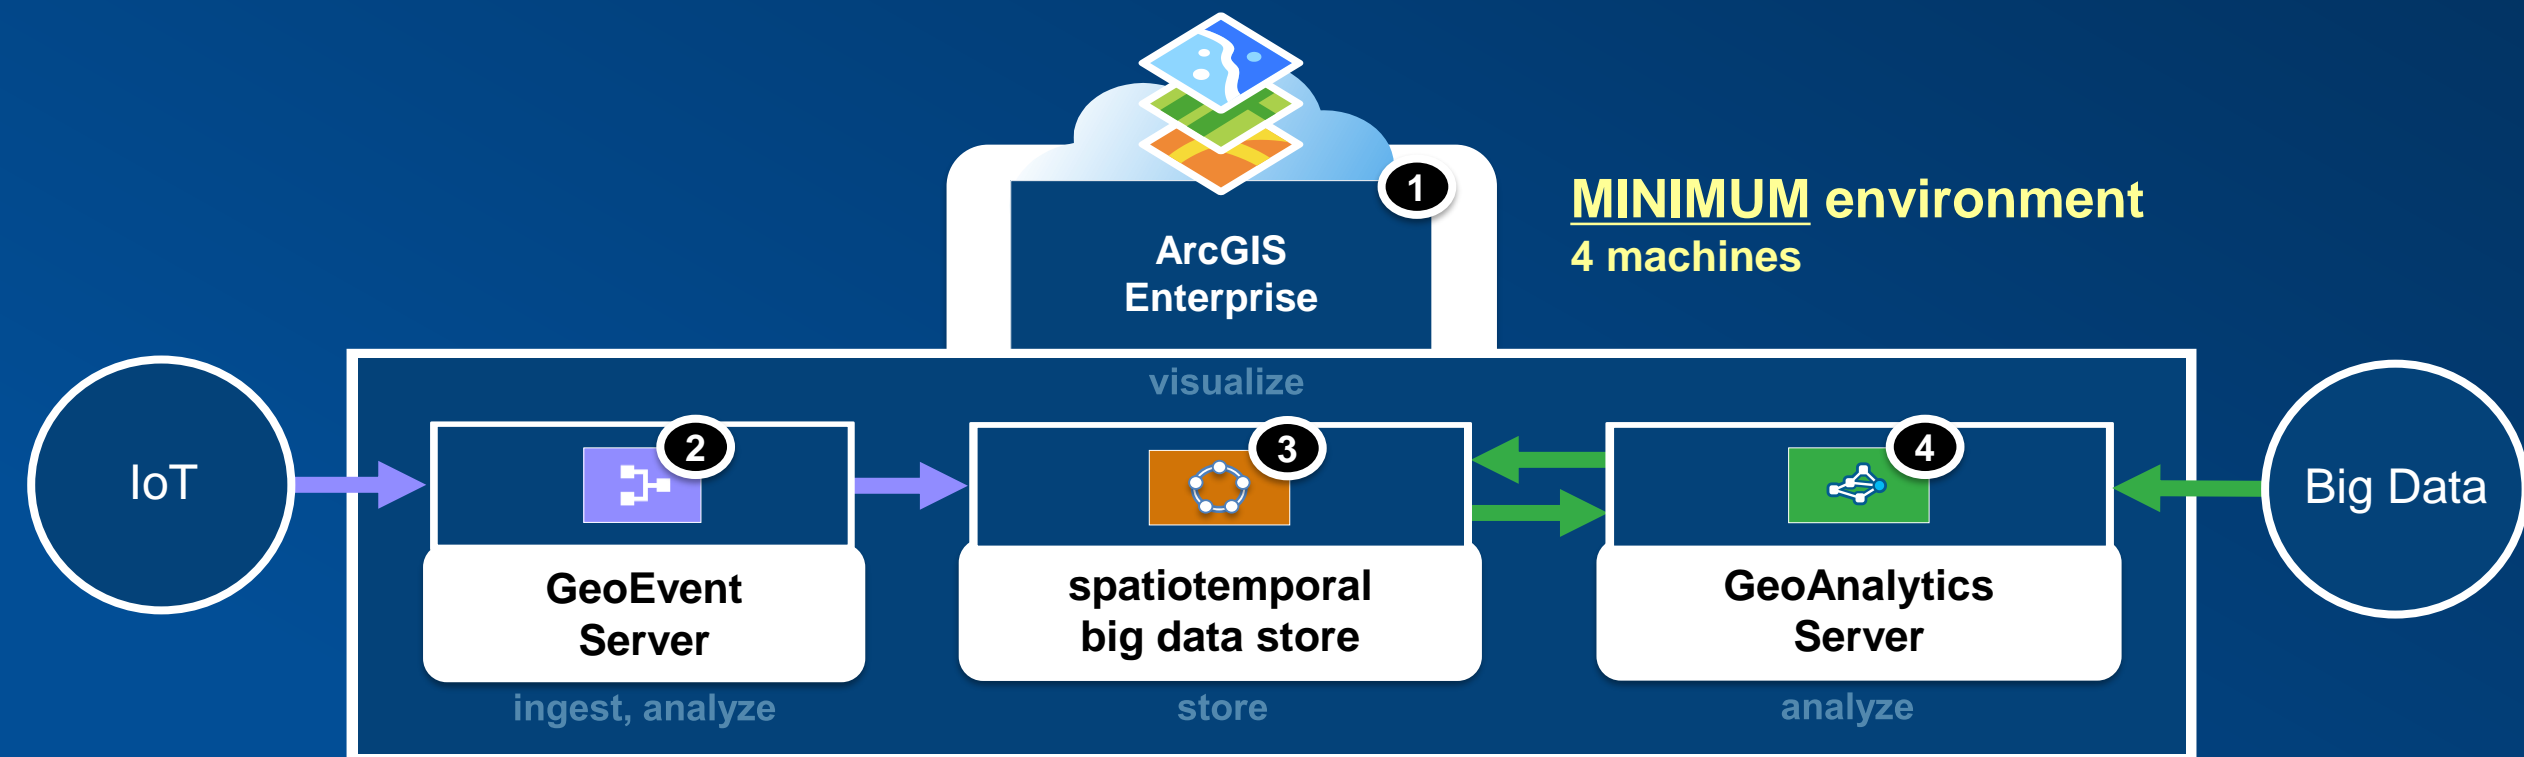
with real-time capabilities



functional servers & spatiotemporal big data store
SHOULD BE on ISOLATED machines

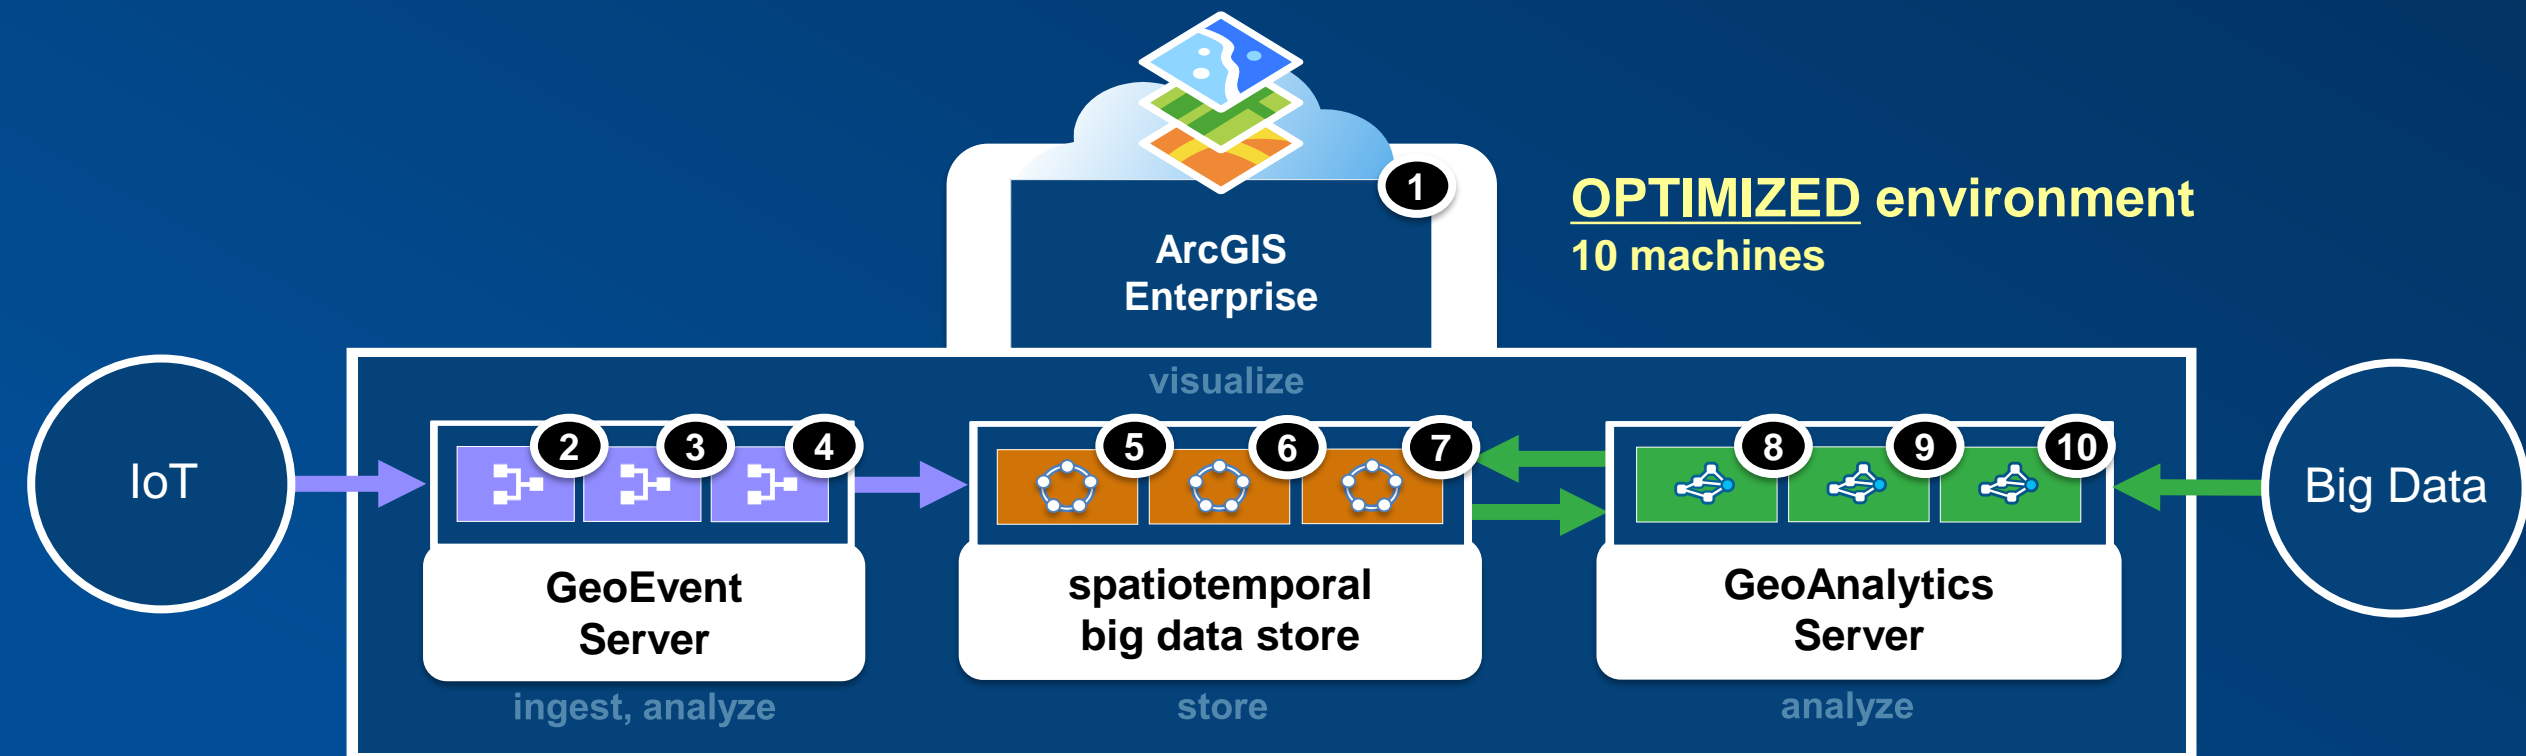
ArcGIS Enterprise

with real-time & big data GIS capabilities



ArcGIS Enterprise

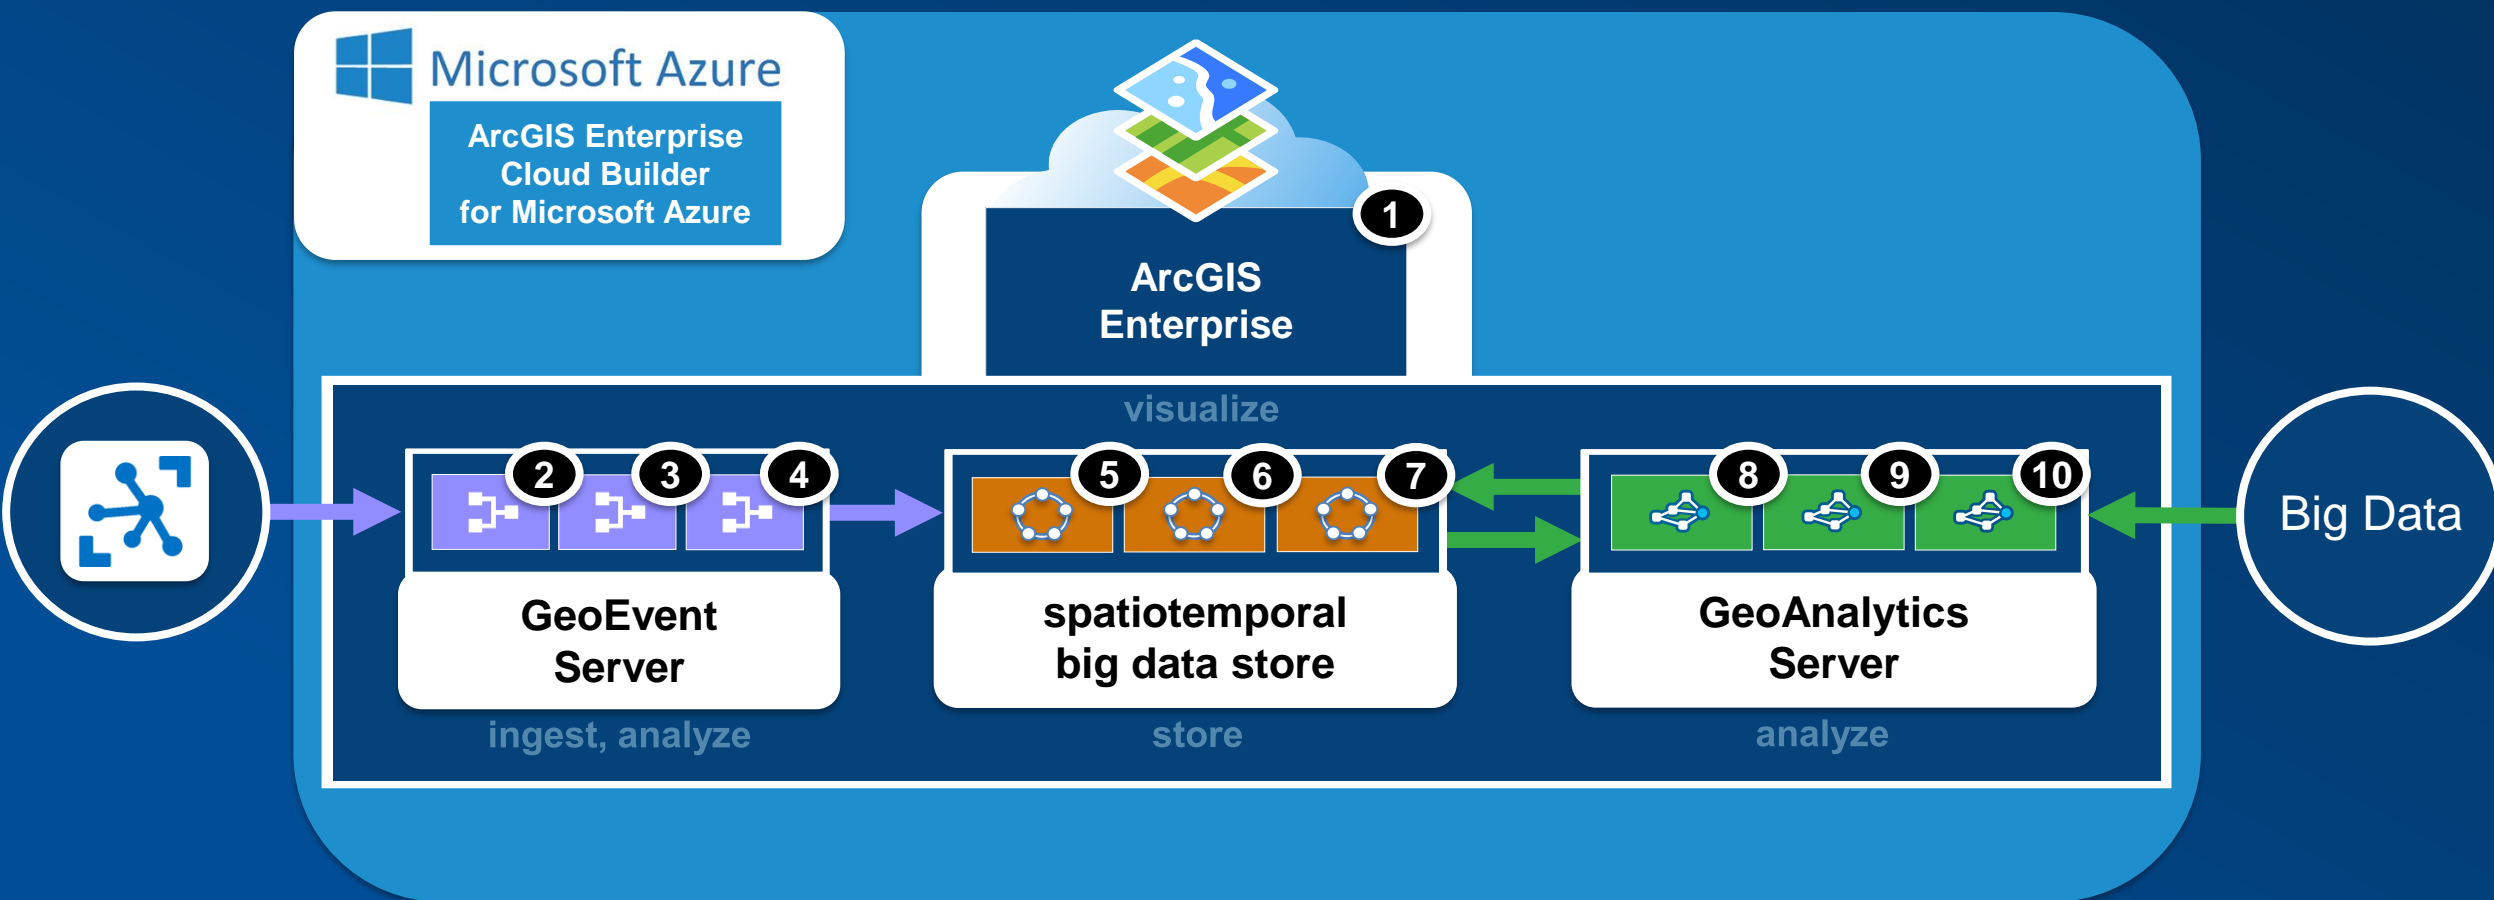
with real-time & big data GIS capabilities



ArcGIS Enterprise

with real-time & big data GIS capabilities on Microsoft Azure

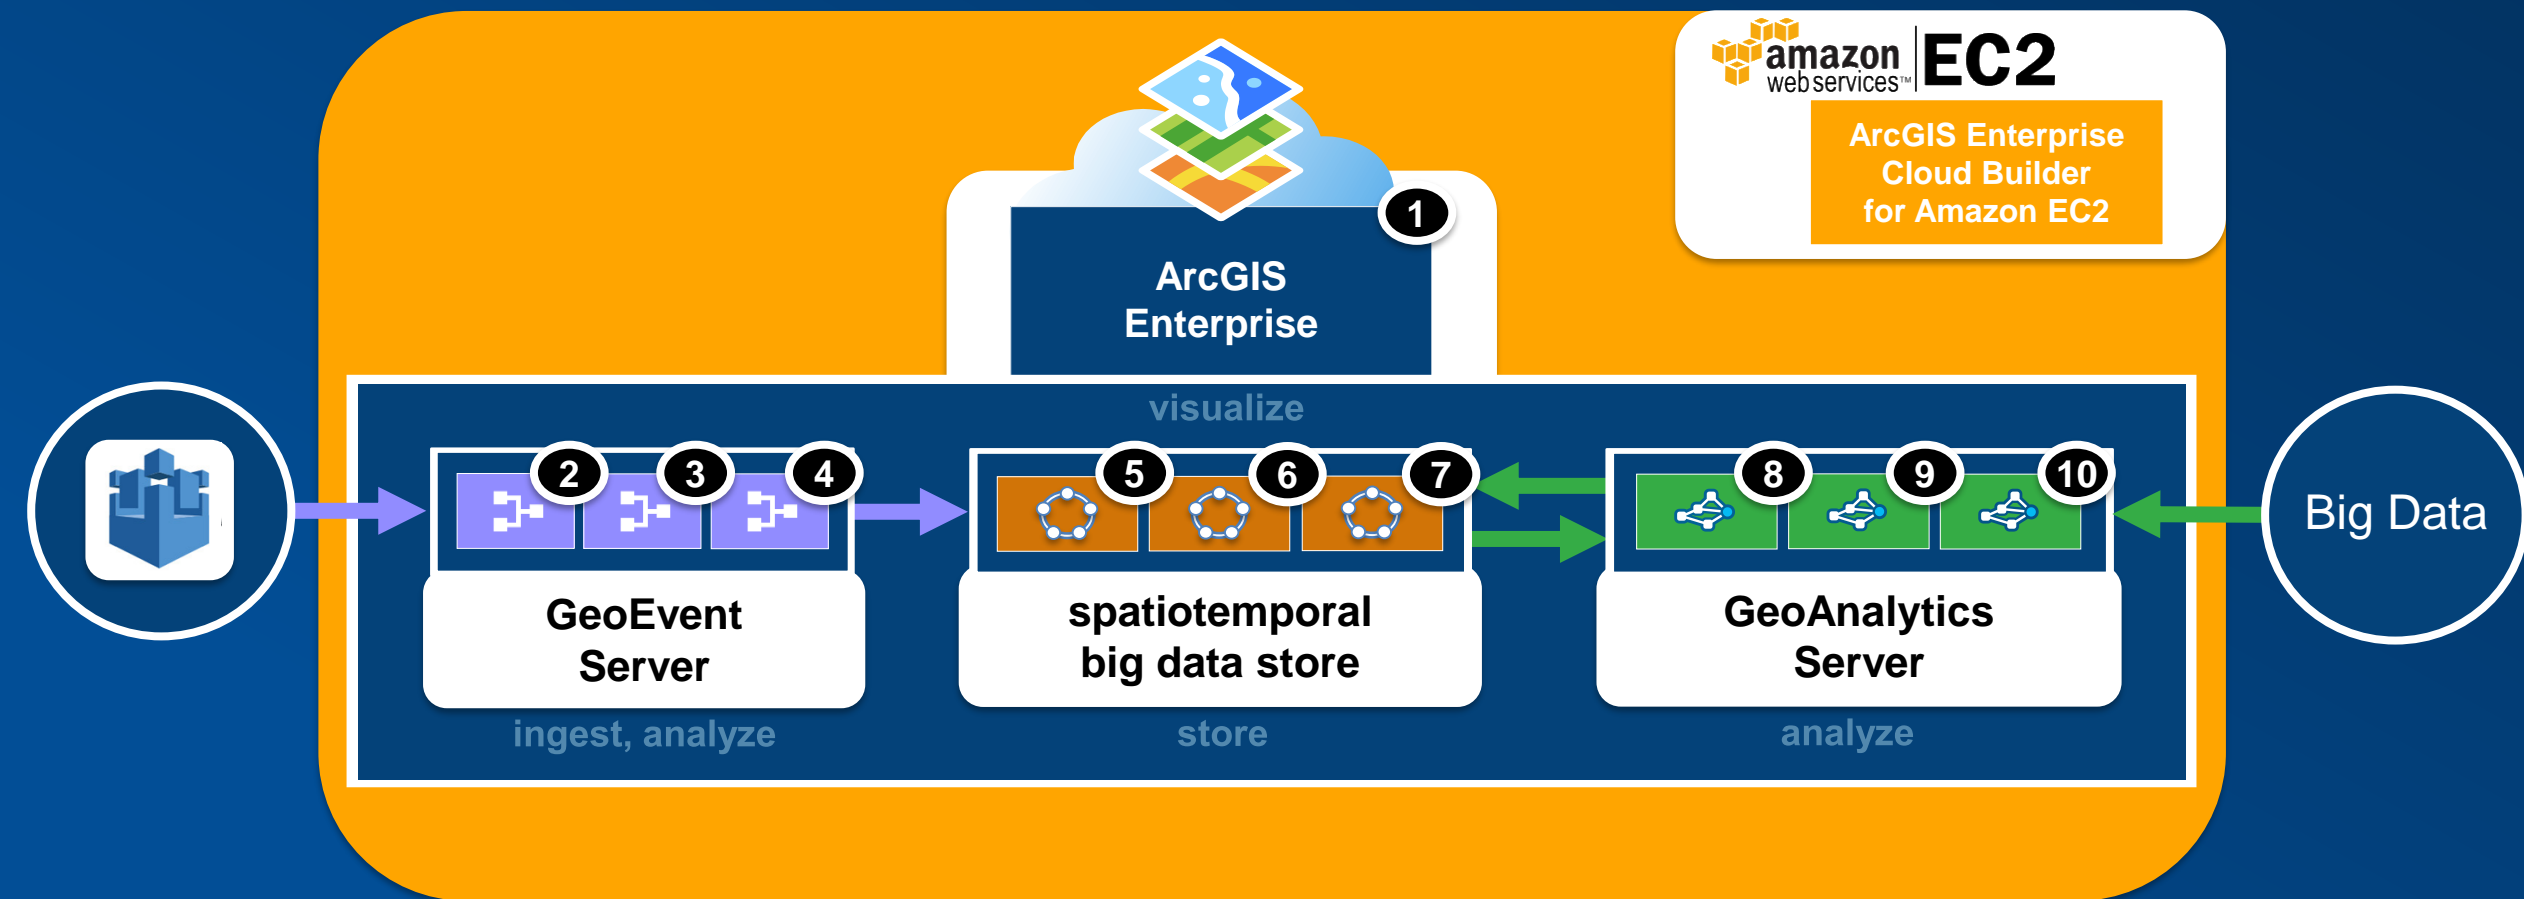
10.6.1



ArcGIS Enterprise

with real-time & big data GIS capabilities on Amazon EC2

10.6.1






Big Data Storage

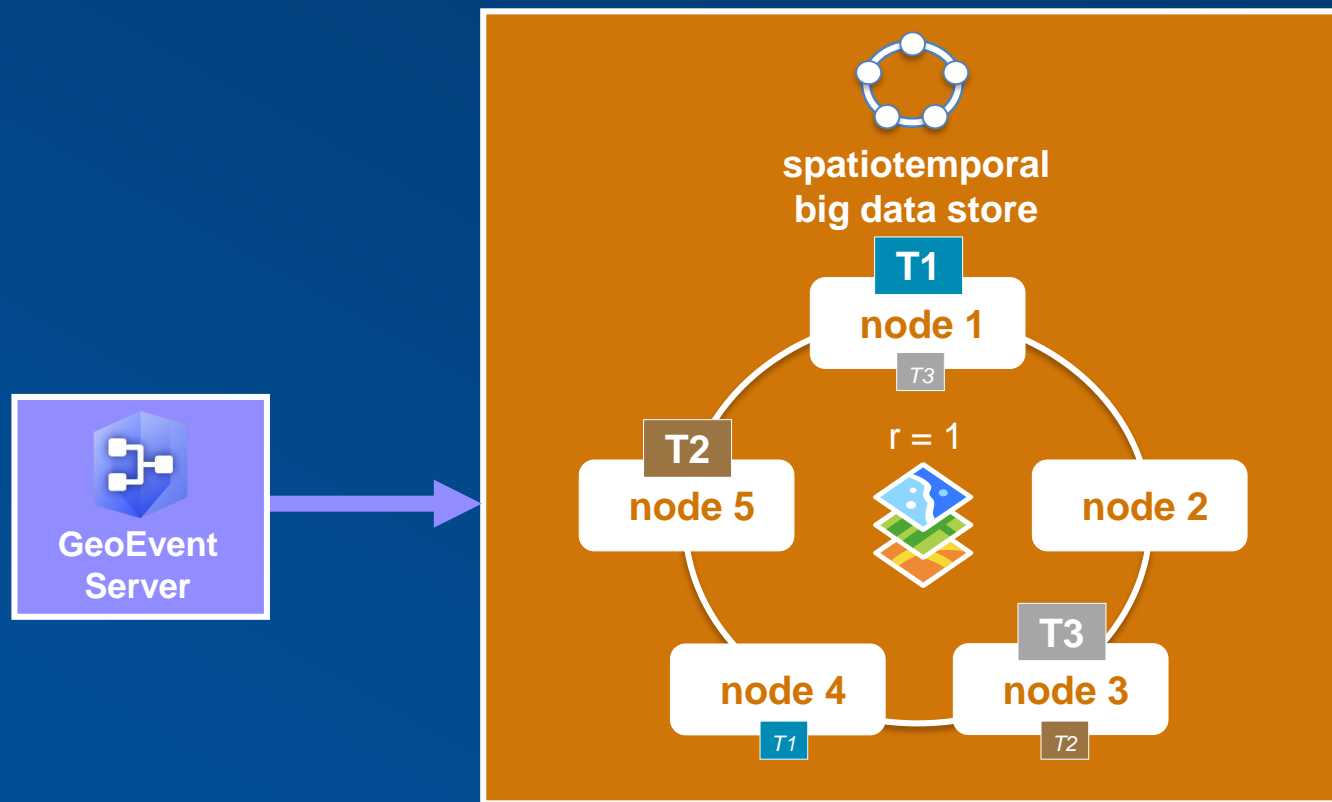
ArcGIS Enterprise

with real-time capabilities



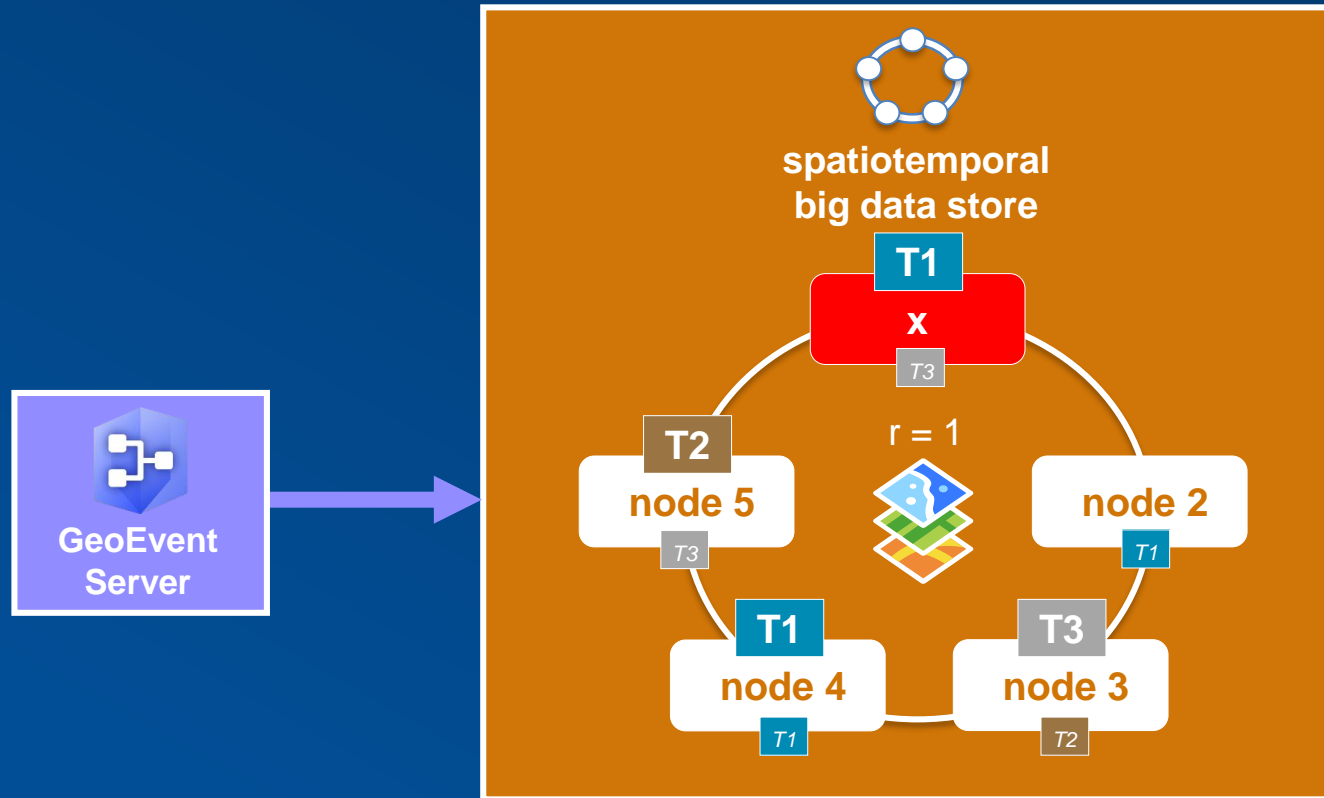
spatiotemporal big data store

shards & replication factor



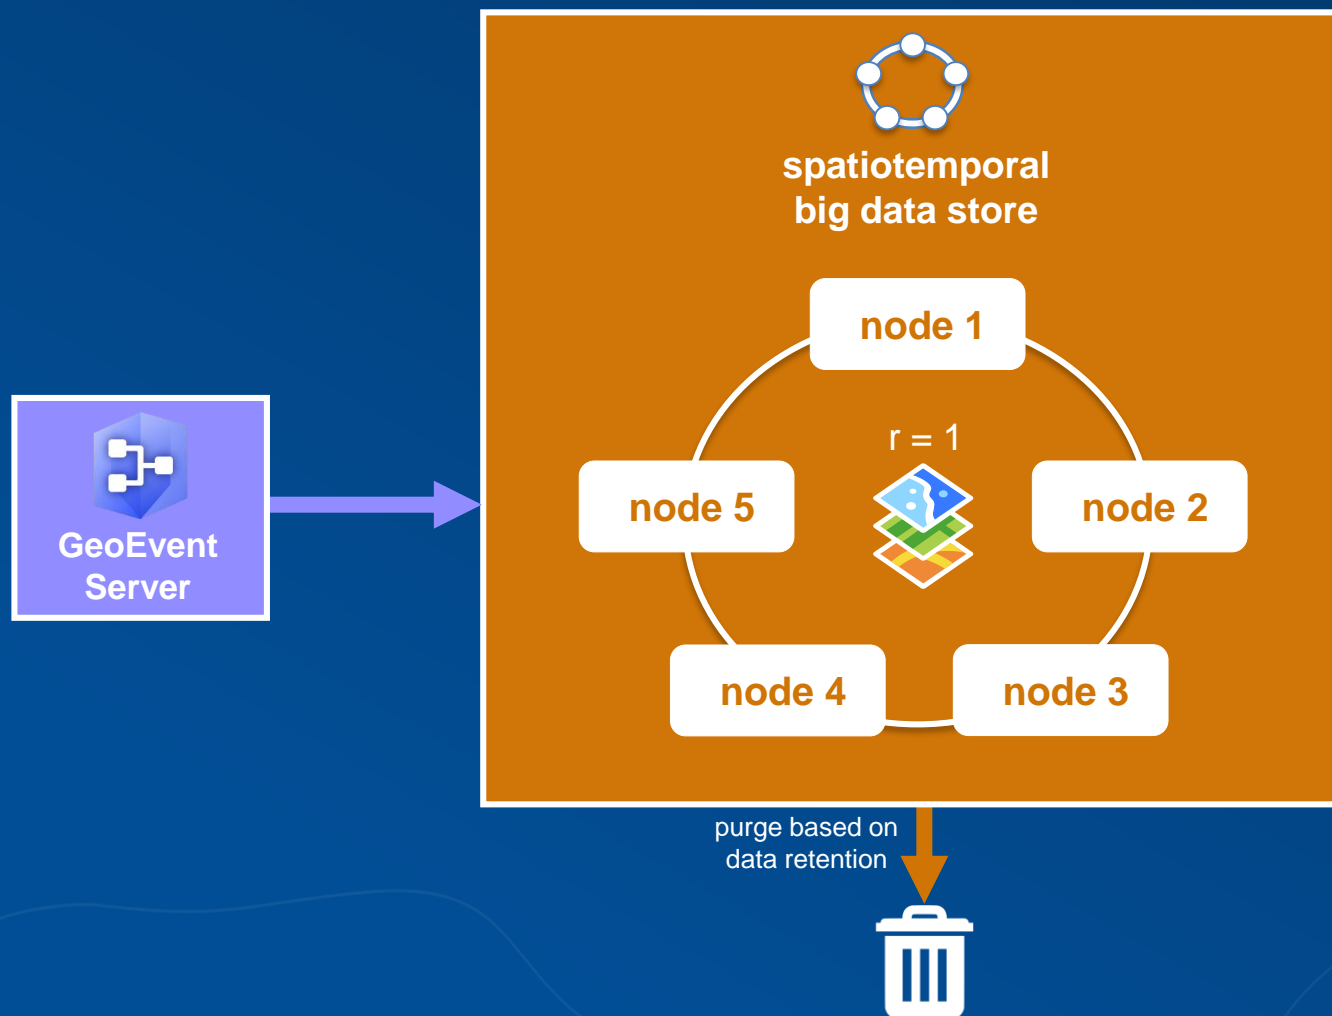
spatiotemporal big data store

auto-rebalancing of data upon node membership changes, + or -, in the big data store



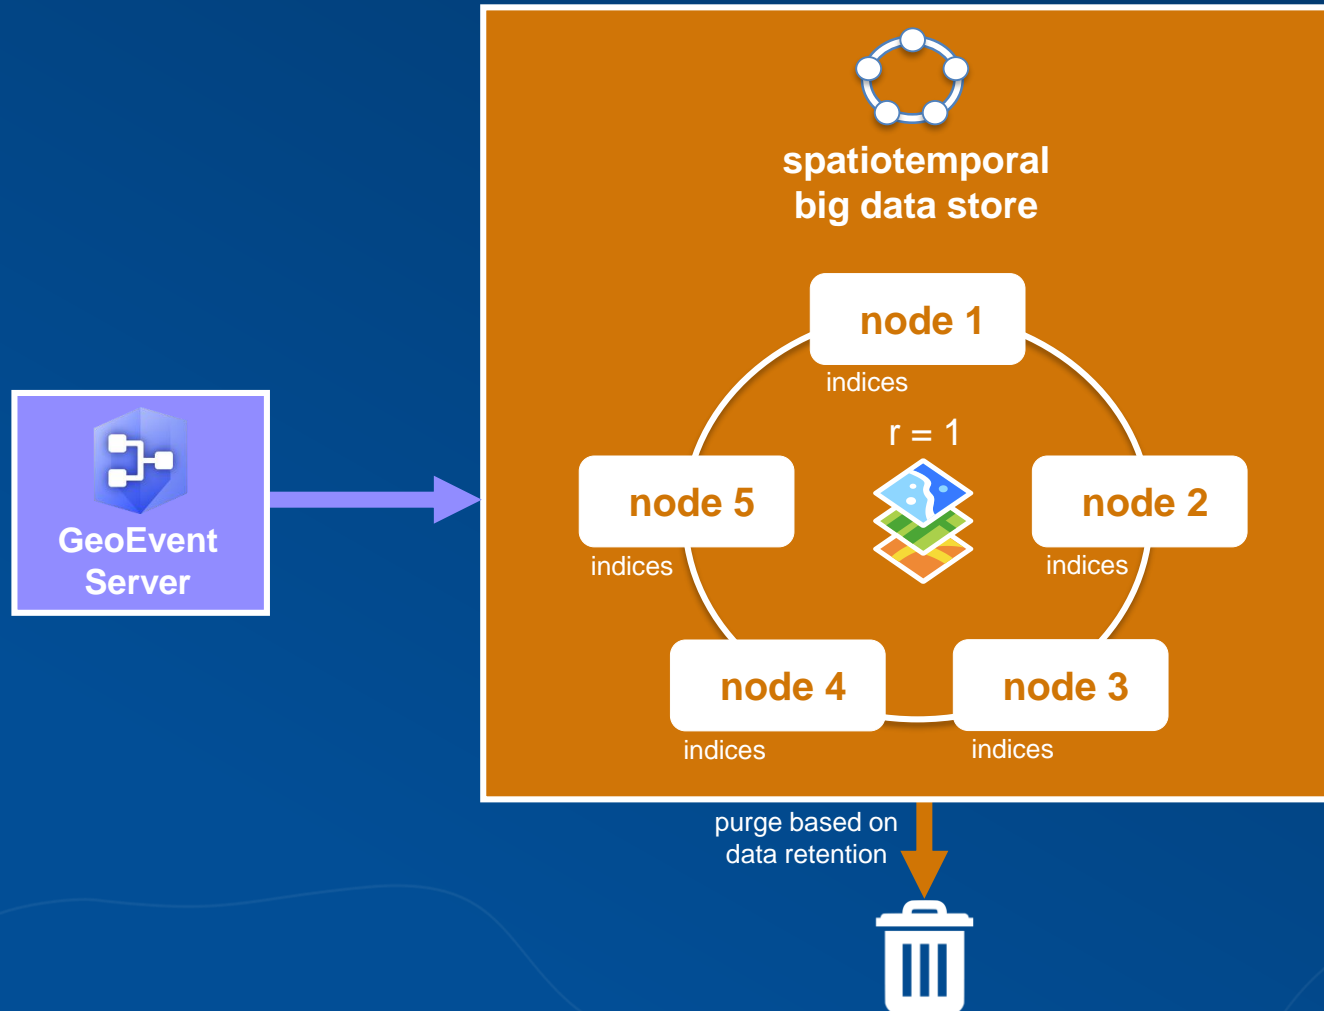
spatiotemporal big data store

data retention policies, configured per data source



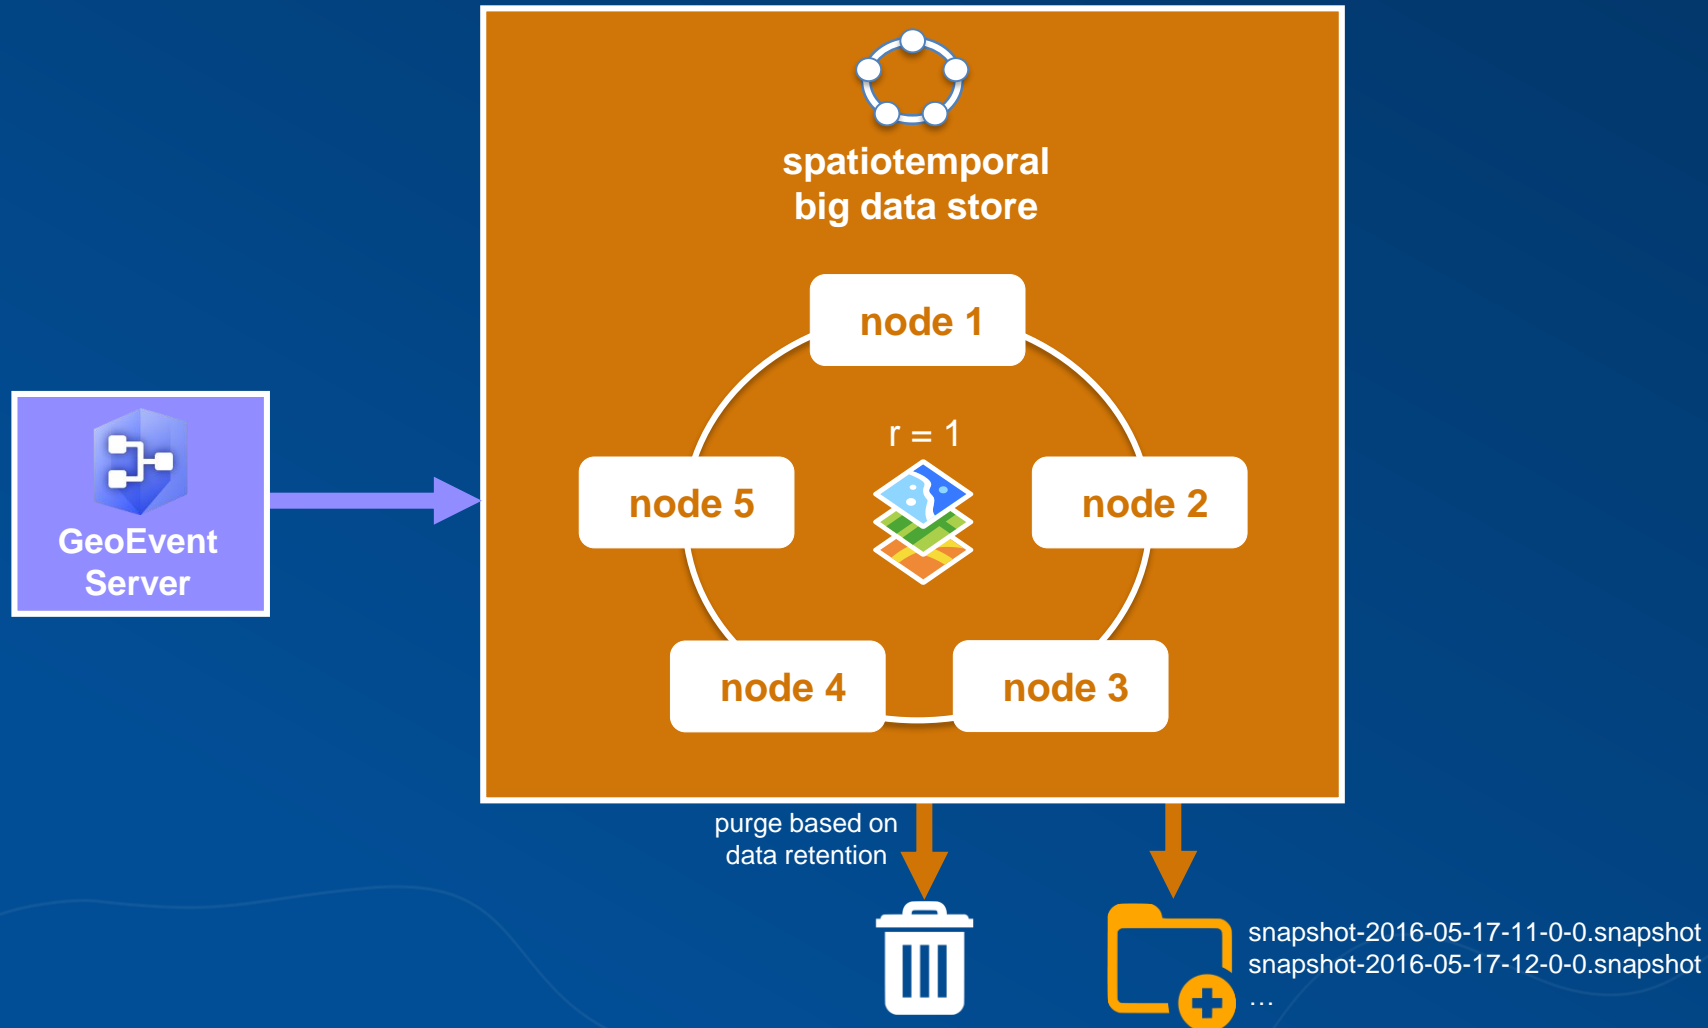
spatiotemporal big data store

rolling index option, set appropriately to the velocity of your observation data



spatiotemporal big data store

automatic data backups using periodic snapshots, including ability to restore from a snapshot



spatiotemporal big data store

choosing an Object Id option

Create Data Source

Name:

GeoEvent Definition:

Geometry Type:

Max Record Count:

▼ Advanced

Replication Factor:

Number of Shards: ☒ Auto

Refresh Interval (seconds):

ObjectID Option:

ObjectID Block Size:

Rolling Data Option:

Data Retention Option: ☐

spatiotemporal big data store

choosing an Object Id option

	Max Value	# of IDs	ArcGIS Clients
Int32	2,147,483,647	2.1 billion	Pro, Desktop, Ops Dashboard, ...

	events per day	Int32
1,000 e/s	86,400,000	25 days
10,000 e/s	864,000,000	2.5 days
100,000 e/s	8,640,000,000	6 hours
1,000,000 e/s	86,400,000,000	36 minutes
10,000,000 e/s	864,000,000,000	4 minutes

spatiotemporal big data store

choosing an Object Id option

	Max Value	# of IDs	ArcGIS Clients
Int32	2,147,483,647	2.1 billion	Pro, Desktop, Ops Dashboard, ...
Int64 (signed)	9,223,372,036,854,775,807	9.2 quintillion	JavaScript, custom apps

	events per day	Int32
1,000 e/s	86,400,000	25 days
10,000 e/s	864,000,000	2.5 days
100,000 e/s	8,640,000,000	6 hours
1,000,000 e/s	86,400,000,000	36 minutes
10,000,000 e/s	864,000,000,000	4 minutes

spatiotemporal big data store

choosing an Object Id option

	Max Value	# of IDs	ArcGIS Clients
Int32	2,147,483,647	2.1 billion	Pro, Desktop, Ops Dashboard, ...
Int64 (signed)	9,223,372,036,854,775,807	9.2 quintillion	JavaScript, custom apps

	events per day	Int32	Int64 (signed)
1,000 e/s	86,400,000	25 days	292,472,000 years
10,000 e/s	864,000,000	2.5 days	29,247,200 years
100,000 e/s	8,640,000,000	6 hours	2,924,720 years
1,000,000 e/s	86,400,000,000	36 minutes	292,472 years
10,000,000 e/s	864,000,000,000	4 minutes	29,248 years

spatiotemporal big data store

choosing an Object Id option

	Max Value	# of IDs	ArcGIS Clients
Int32	2,147,483,647	2.1 billion	Pro, Desktop, Ops Dashboard, ...
Int64 (signed)	9,223,372,036,854,775,807	9.2 quintillion	JavaScript, custom apps
UniqueStringID	n/a	unlimited	JavaScript, custom apps

	events per day	Int32	Int64 (signed)
1,000 e/s	86,400,000	25 days	292,472,000 years
10,000 e/s	864,000,000	2.5 days	29,247,200 years
100,000 e/s	8,640,000,000	6 hours	2,924,720 years
1,000,000 e/s	86,400,000,000	36 minutes	292,472 years
10,000,000 e/s	864,000,000,000	4 minutes	29,248 years



Performance, Resiliency & Scalability

Performance, Resiliency & Scalability

throughput benchmarks

- Throughput has increased with every release of ArcGIS GeoEvent Server
 - 10.6 release can support up to 6,000 events per second (e/s)
- Testing harness and sample data will be made available on GitHub
- Check our blog on GeoNet after the conference for details next week

<http://links.esri.com/geoevent-forum>

ArcGIS GeoEvent Server	10.2	10.3	10.4	10.5	10.6
Velocity throughput <i>measured in events per second (e/s)</i>	<i>up to 500 e/s</i>	<i>up to 2,000 e/s</i>	<i>up to 3,000 e/s</i>	<i>up to 4,000 e/s</i>	<i>up to 6,000 e/s</i>

Performance, Resiliency & Scalability

multi-machine site support

ArcGIS 10.5

- Resiliency (high availability) & scalability is only possible if users “bring their own gateway”.
 - Barrier to entry is HIGH & typically requires a professional services engagement for success.
- Loses flexibility of input types.

ArcGIS 10.6

- Provides users with a resilient & scalable Real-Time GIS deployment OUT-OF-THE-BOX.
 - Introduces a gateway process that is automatically configured as part of GeoEvent Server installation.
- Provides flexibility for all input types.

ArcGIS GeoEvent Server	10.2	10.3	10.4	10.5	10.6
Velocity throughput <i>measured in events per second (e/s)</i>	<i>up to 500 e/s</i>	<i>up to 2,000 e/s</i>	<i>up to 3,000 e/s</i>	<i>up to 4,000 e/s</i>	<i>up to 6,000 e/s</i>
Resiliency & Scalability <i>via multi-machine site</i>	no	no	no	no	yes

Performance, Resiliency & Scalability

multi-machine site support

ArcGIS GeoEvent Server – Multiple-Machine Site Tutorial available online now

Tutorial - GeoEvent Server 10.6.x Multiple-Machine Site

Overview



ArcGIS GeoEvent Server 10.6.x now supports the creation of multiple-machine sites.

Document Link by [GeoEventTeam](#)

Created: Mar 3, 2018 Updated: Mar 3, 2018

View Count: 0

Open

Description

ArcGIS GeoEvent Server 10.6.x now supports the creation of multiple-machine sites. In a multiple-machine site, two or more GeoEvent Server machines can be administered and used as a single logical unit, providing GeoEvent Server administrators with great flexibility to easily adjust the computing power of the site by adding or removing GeoEvent Server machines.

This tutorial will walk you through how to plan, setup, and work with a GeoEvent Server 10.6.x multiple-machine site. Also included is an appendix for administrators to learn how to monitor an existing GeoEvent Server multiple-machine site.

Details

Size: 1 KB

★★★★★


Owner


 GeoEventTeam

Tags

[arcgis](#), [geoevent](#), [server](#), [real](#), [time](#), [real-time](#), [realtime](#), [multiple](#), [machine](#), [site](#), [tutorial](#), [scale](#), [scaling](#),

ArcGIS® GeoEvent Server Multiple-Machine Site Tutorial



 **esri** | THE SCIENCE OF WHERE™

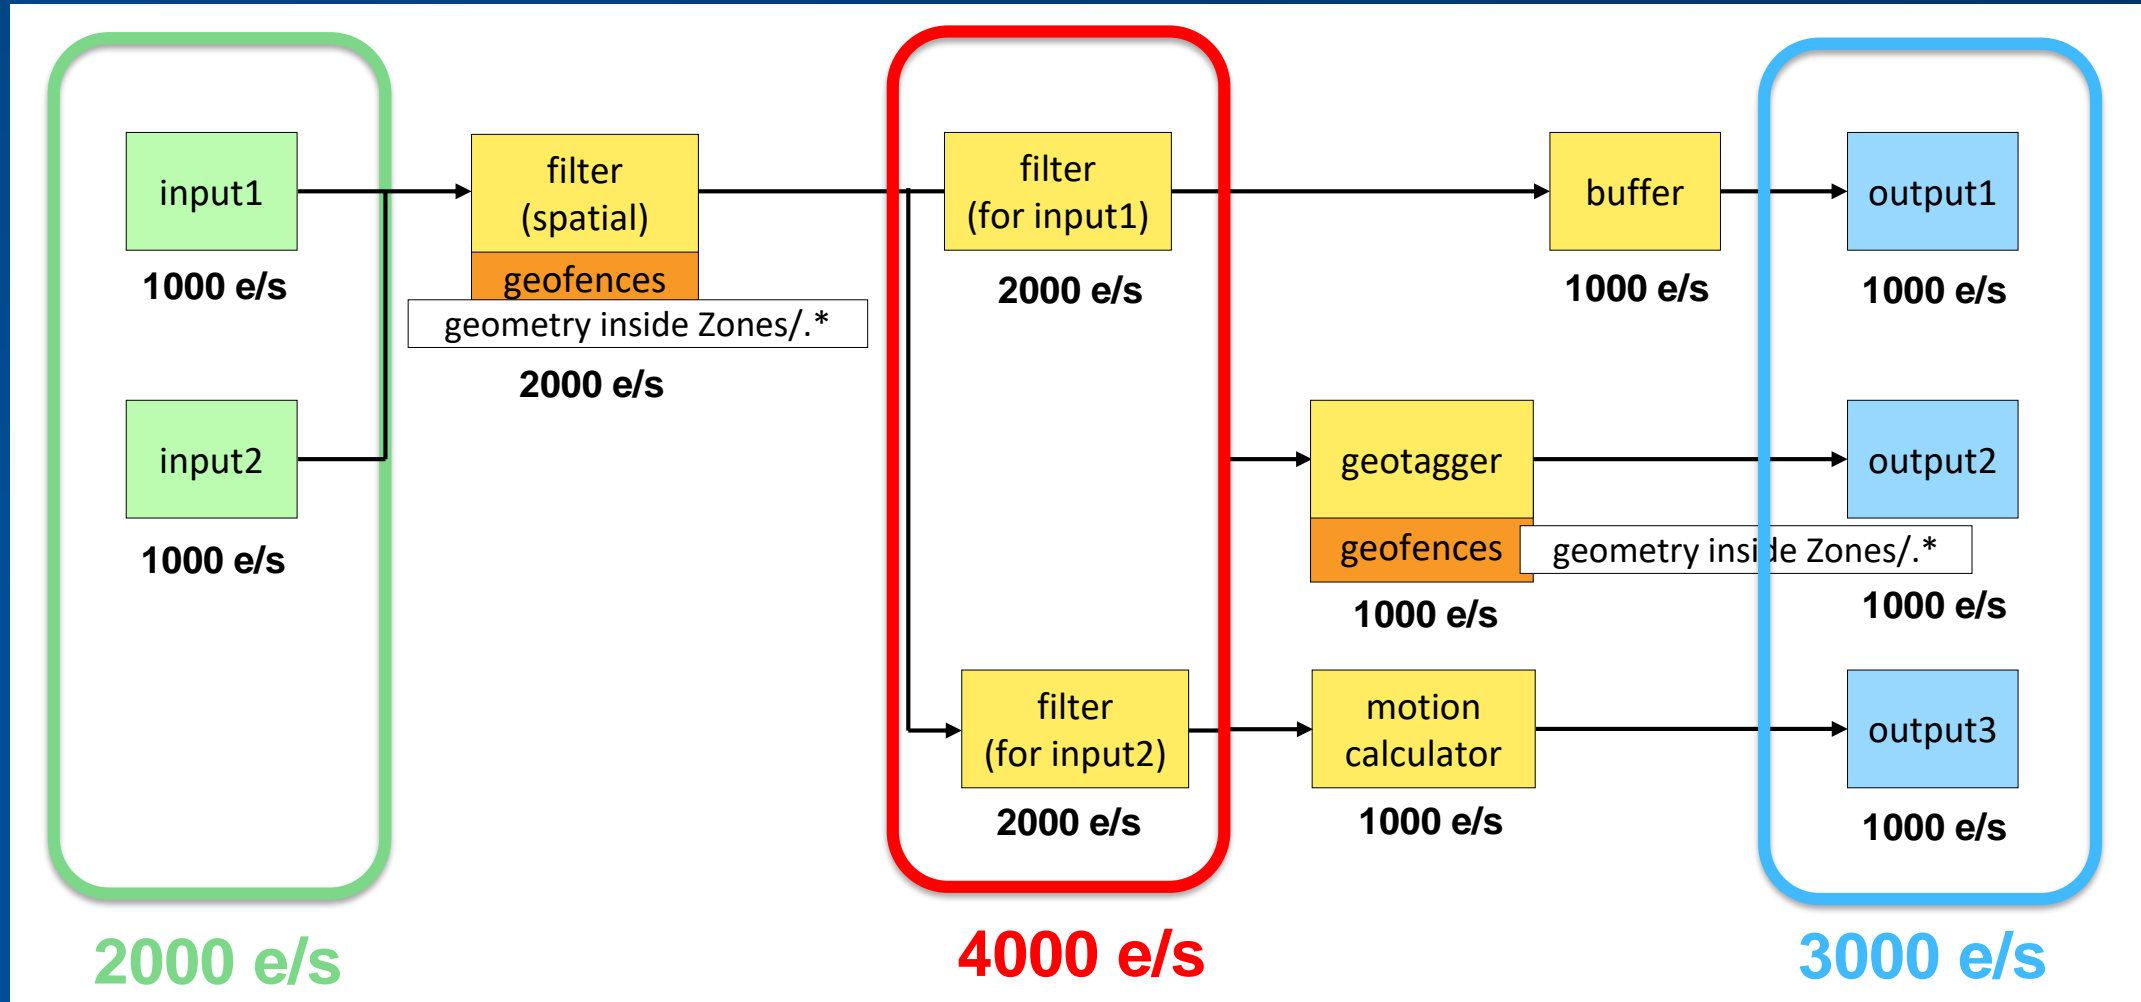
NOTE: The GeoEvent Server Team strives to update product tutorials to reflect the latest release. Depending on the version of GeoEvent Server you are using, there may be inconsistencies between your environment and the illustrations and/or specific steps in exercises or videos bundled with the tutorial. The concepts outlined, however, should be applicable across different versions of GeoEvent Server.

Tutorial - GeoEvent Server Multiple-Machine Site 10.6.x / r1

<http://links.esri.com/geoevent-multiplemachine>

Performance, Resiliency & Scalability

factors that influence throughput



Input event counts don't always tell the whole story

Performance, Resiliency & Scalability

configuration changes to support larger scale

- GeoEvent Server by default is only allocated 4GB of RAM for the JVM
 - If utilizing a large amount of GeoFences it may be necessary to increase this amount
- This can be modified through the “/etc/ArcGISGeoEvent.cfg” up to 32GB (JVM limitation)

```
wrapper.java.additional.12=-XX:+CMSClassUnloadingEnabled
```

```
# Maximum Java Heap Size (in MB). Note: use this option verse the one below 'wrapper.java.maxmemory' when setting the max heap size.  
wrapper.java.additional.13=-Xmx4096m
```

```
# Optional Parameter to control the max file size in run configuration store (zookeeper). 10xMB = 1048576 bytes = 10 MB. The default is 10 MB.
```

```
# Initial Java Heap Size (in MB)  
#wrapper.java.initmemory=128
```





Performance, Resiliency & Scalability

configuration changes to support larger scale

- GeoEvent Server by default is able to maintain the state of 1000 unique Track_IDs.
 - This value can be increased by editing “/etc/com.esri.ges.manager.servicemanager.cfg”

```
com.esri.ges.manager.servicemanager.maxCacheSize=1000
```

- You may also need to modify the Incident Manager Setting in the Global Setting Tab if used in conjunction with the Incident Detector Processor

▾ Incident Manager Settings		
Maximum number of closed incidents	1000	 
Maximum number of opened incidents	1000	 



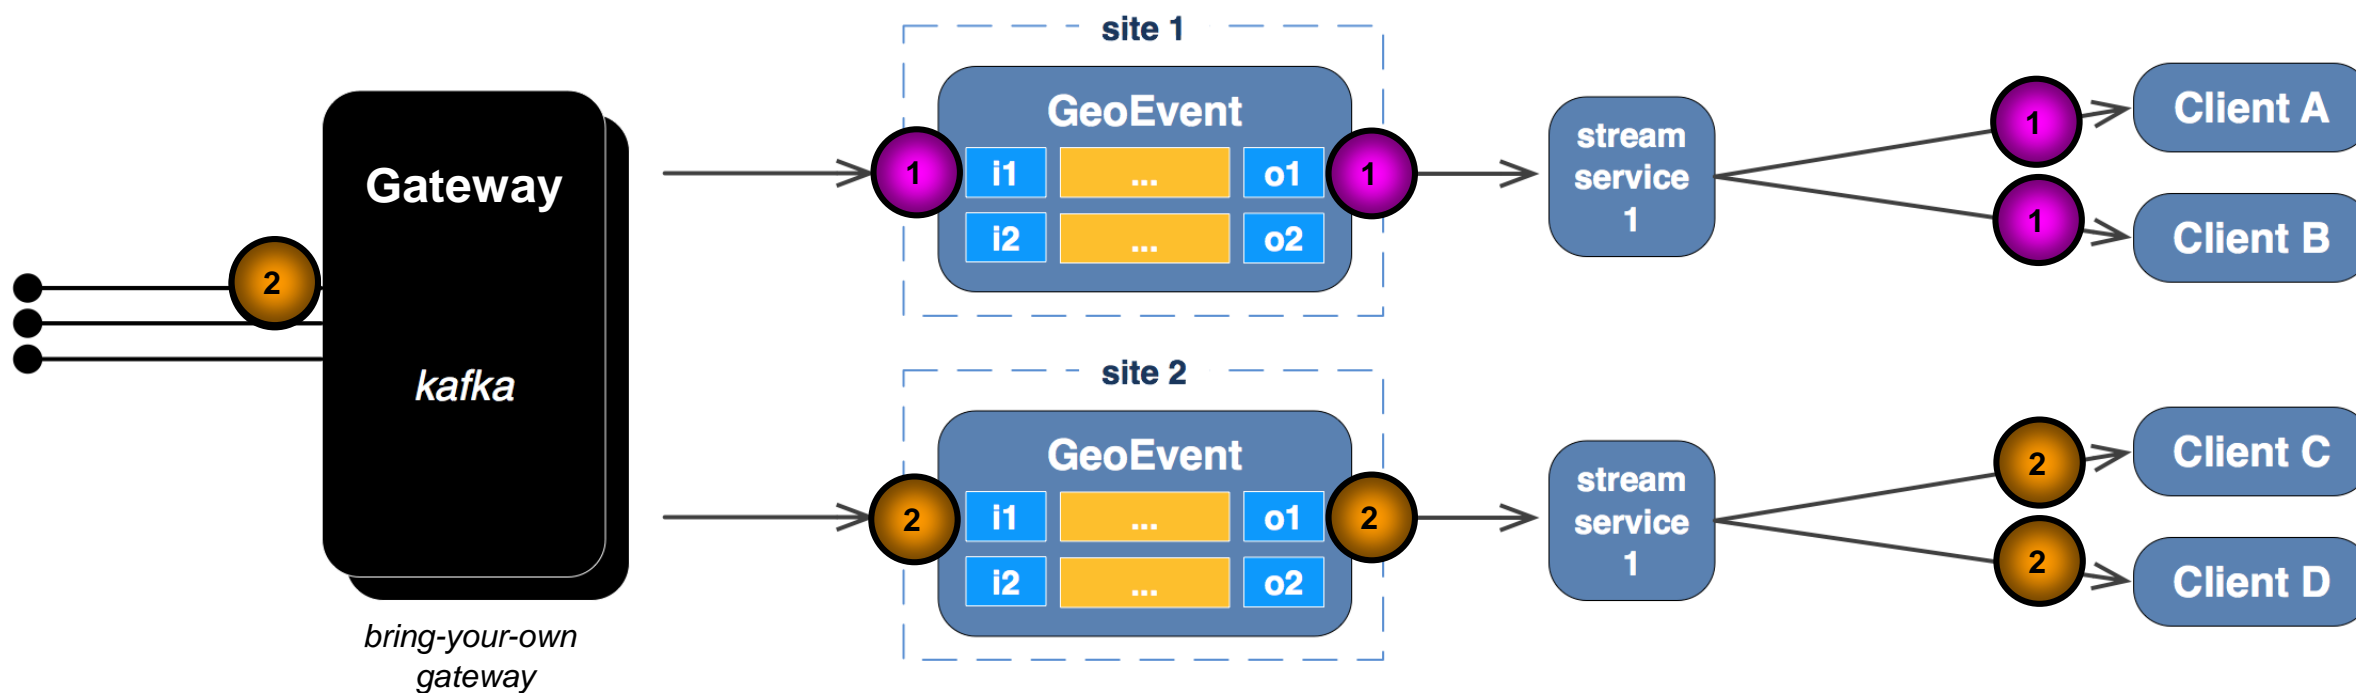
Stream Services

Stream Service resilience & scalability

10.5 best practice = isolated deployment of GeoEvent Server (site per GeoEvent)

10.5

- An isolated deployment of GeoEvent instances leads to challenges with Stream Services:
 - Client A & B see event 1, while client C & D do not
 - Client C & D see event 2, while client A & B do not

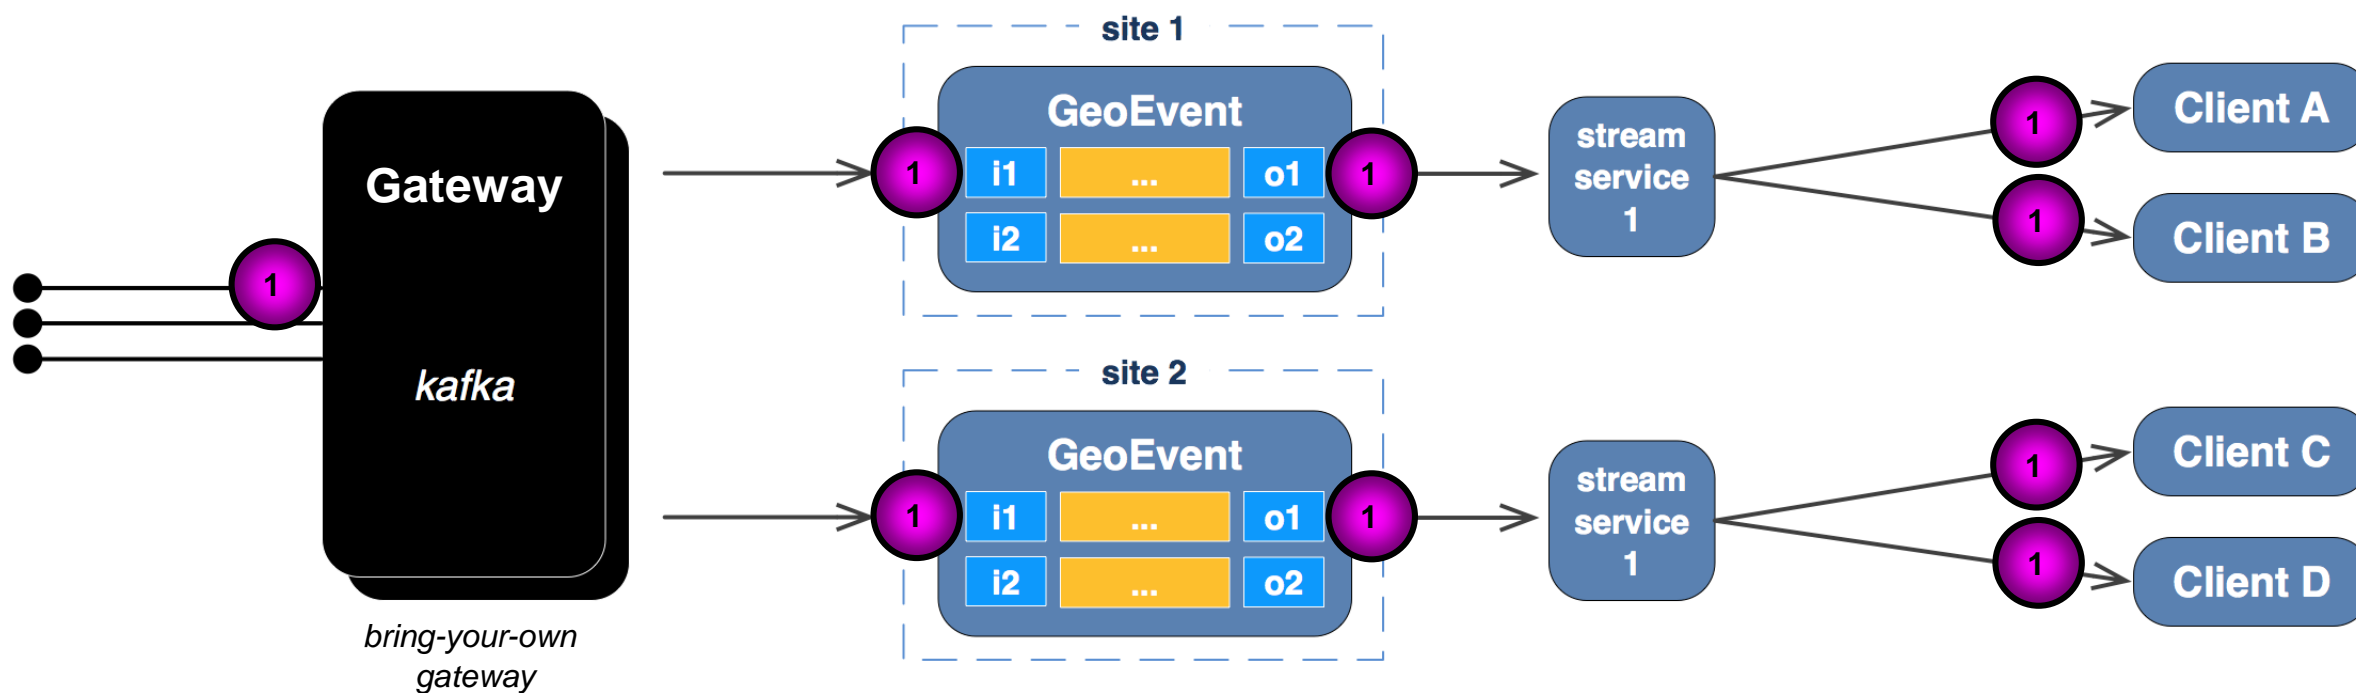


Stream Service resilience & scalability

10.5 best practice = isolated deployment of GeoEvent Server (site per GeoEvent)

10.5

- Using Kafka the GeoEvent instances can be configured to use separate consumer groups:
 - With this configuration, all clients see all events

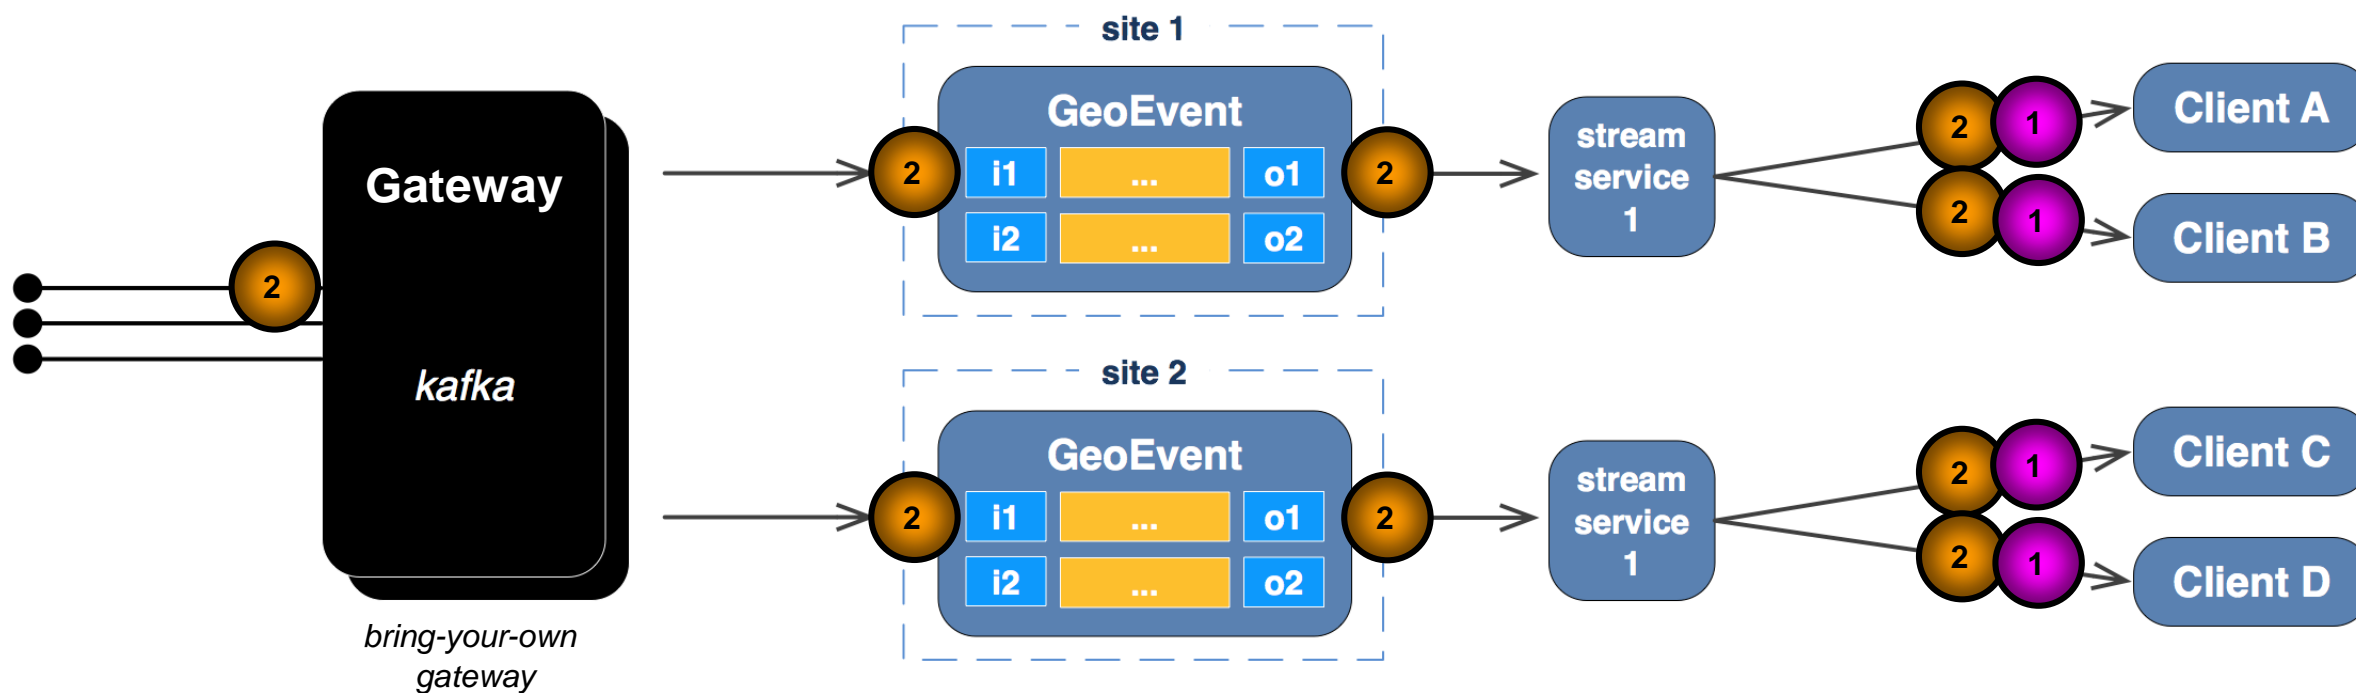


Stream Service resilience & scalability

10.5 best practice = isolated deployment of GeoEvent Server (site per GeoEvent)

10.5

- Using Kafka the GeoEvent instances can be configured to use separate consumer groups:
 - With this configuration, all clients see all events

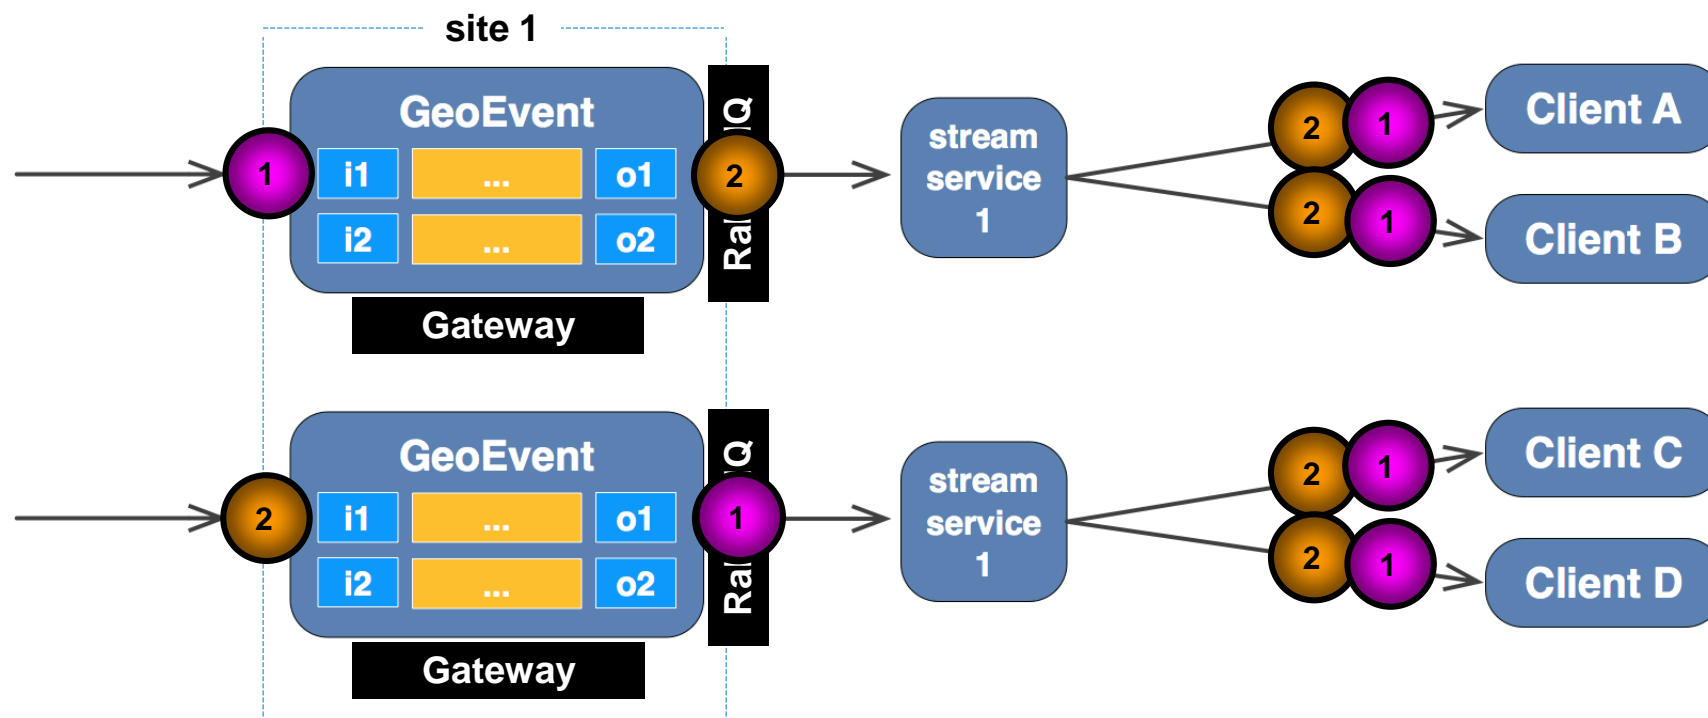


Stream Service resilience & scalability

10.6 best practice = multi-machine site of GeoEvent Servers

10.6

- Gateway is provided out-of-the-box at 10.6 for ingress:
 - all clients see all events by default

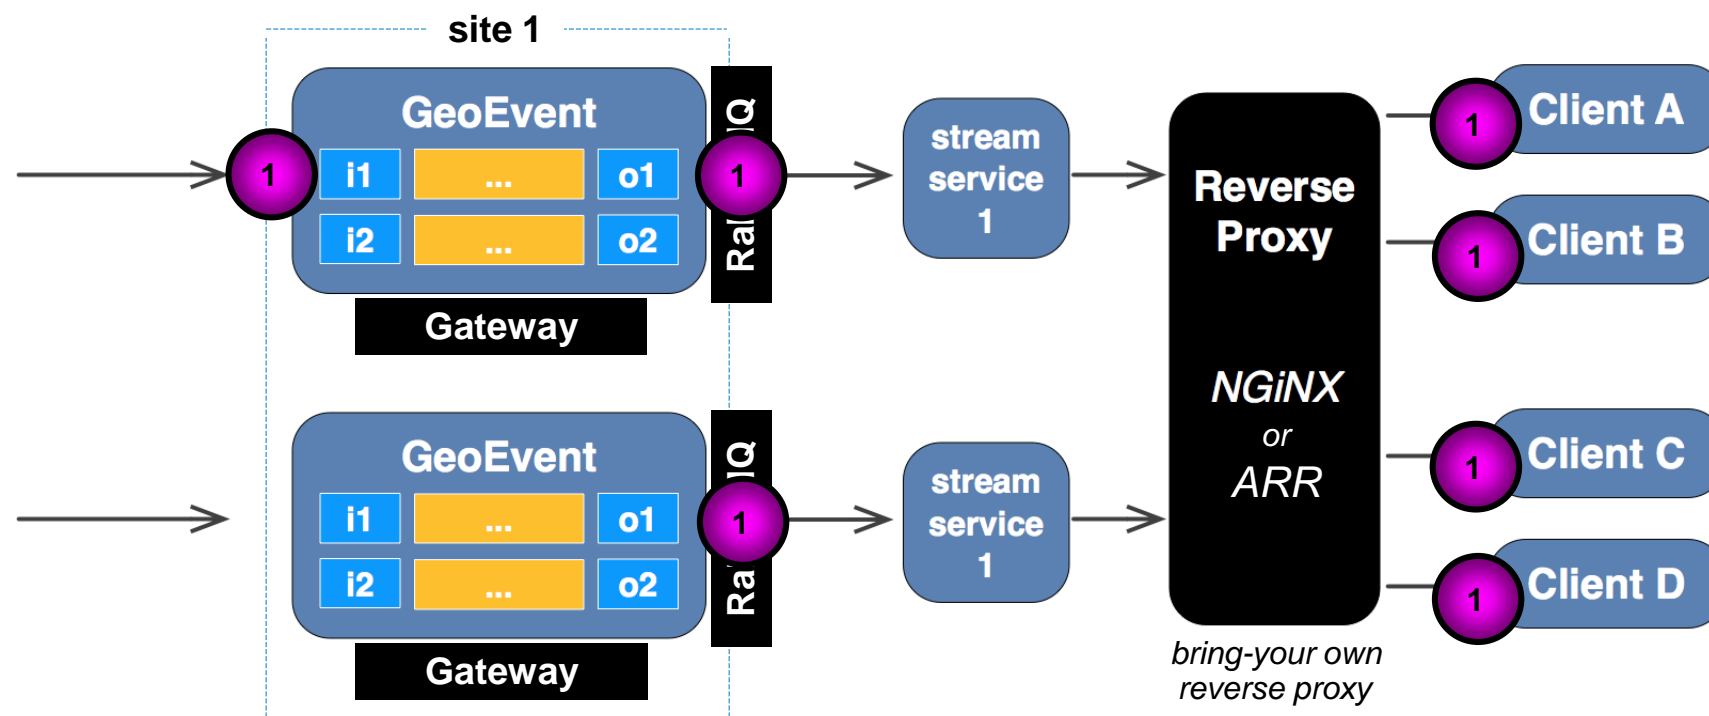


Stream Service resilience & scalability

10.6 best practice = multi-machine site of GeoEvent Servers

10.6

- A reverse proxy can be configured in between the clients and the stream services so that clients don't have direct knowledge of the servers they are connecting to.
 - Example reverse proxies include NGiNX & Microsoft Application Request Routing (ARR).

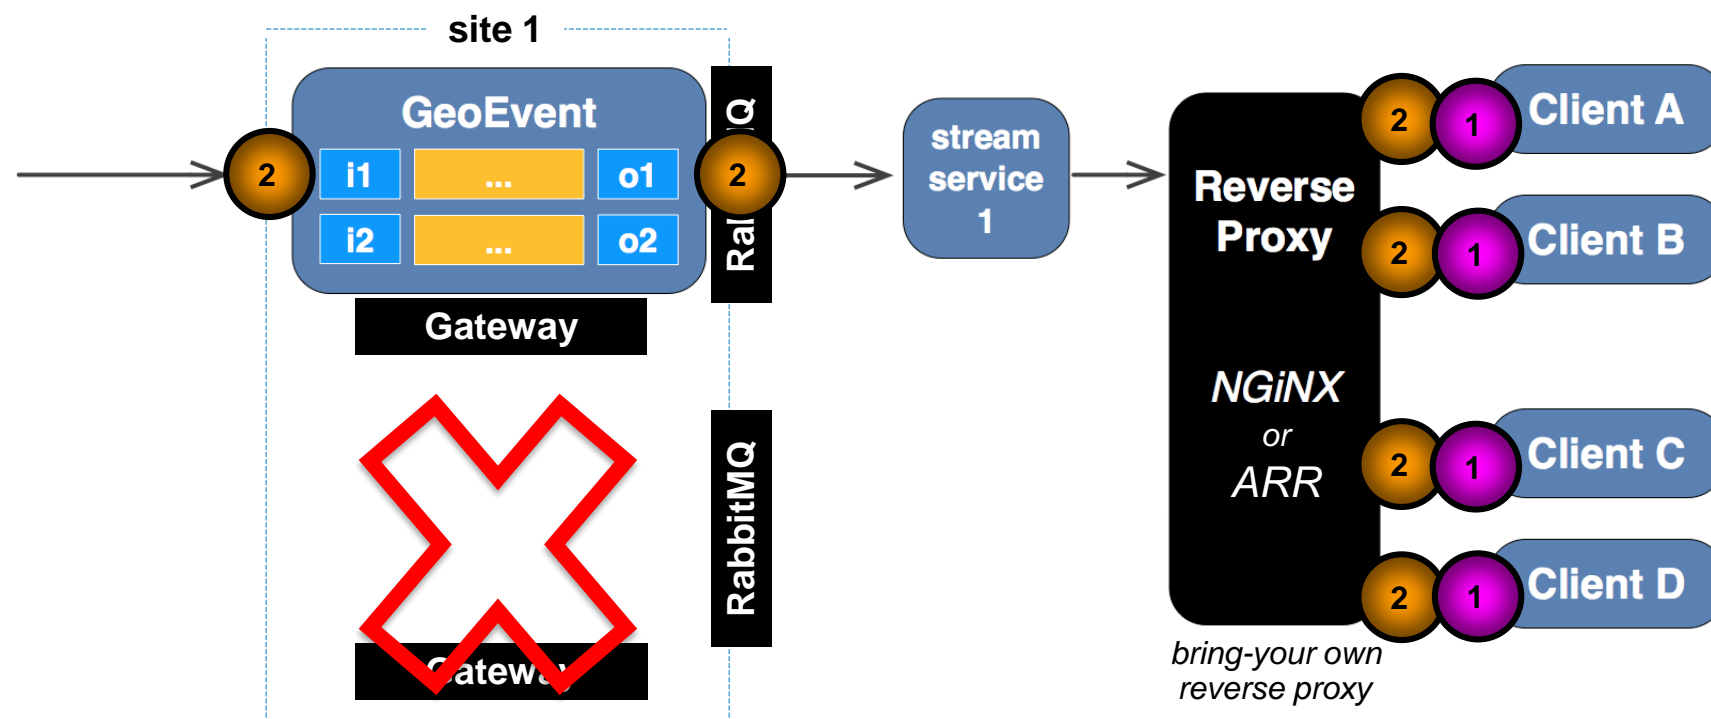


Stream Service resilience & scalability

10.6 best practice = multi-machine site of GeoEvent Servers

10.6

- A reverse proxy can be configured in between the clients and the stream services so that clients don't have direct knowledge of the servers they are connecting to.
 - Example reverse proxies include NGiNX & Microsoft Application Request Routing (ARR).

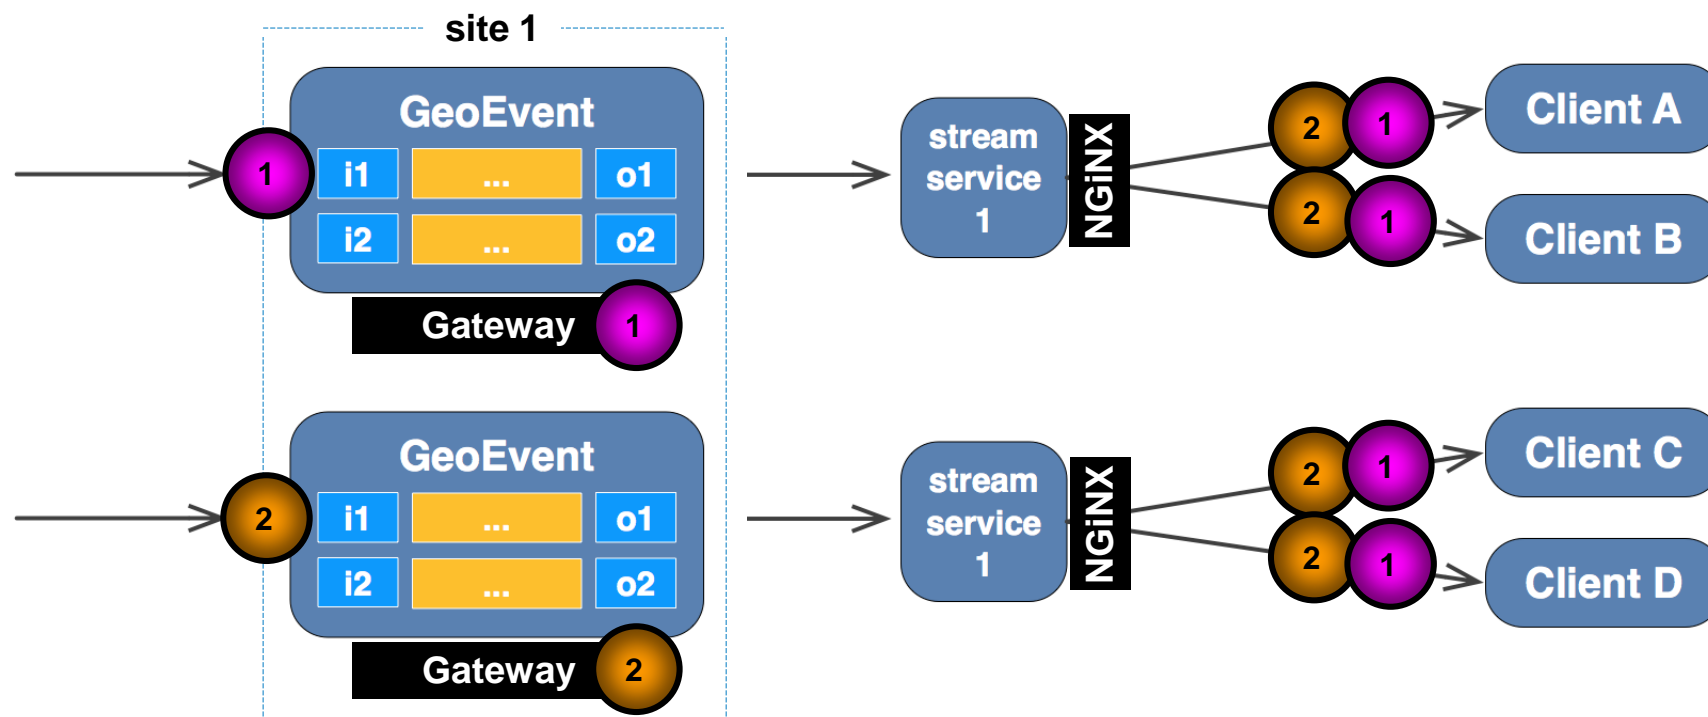


Stream Service resilience & scalability

10.7 best practice = multi-machine site of GeoEvent Servers

10.7

- Gateway usage is expanded to be used by Stream Service egress at 10.7:
 - all clients see all events by default in a much more performant way







Service Design Considerations

Service Design Considerations

which would you choose?



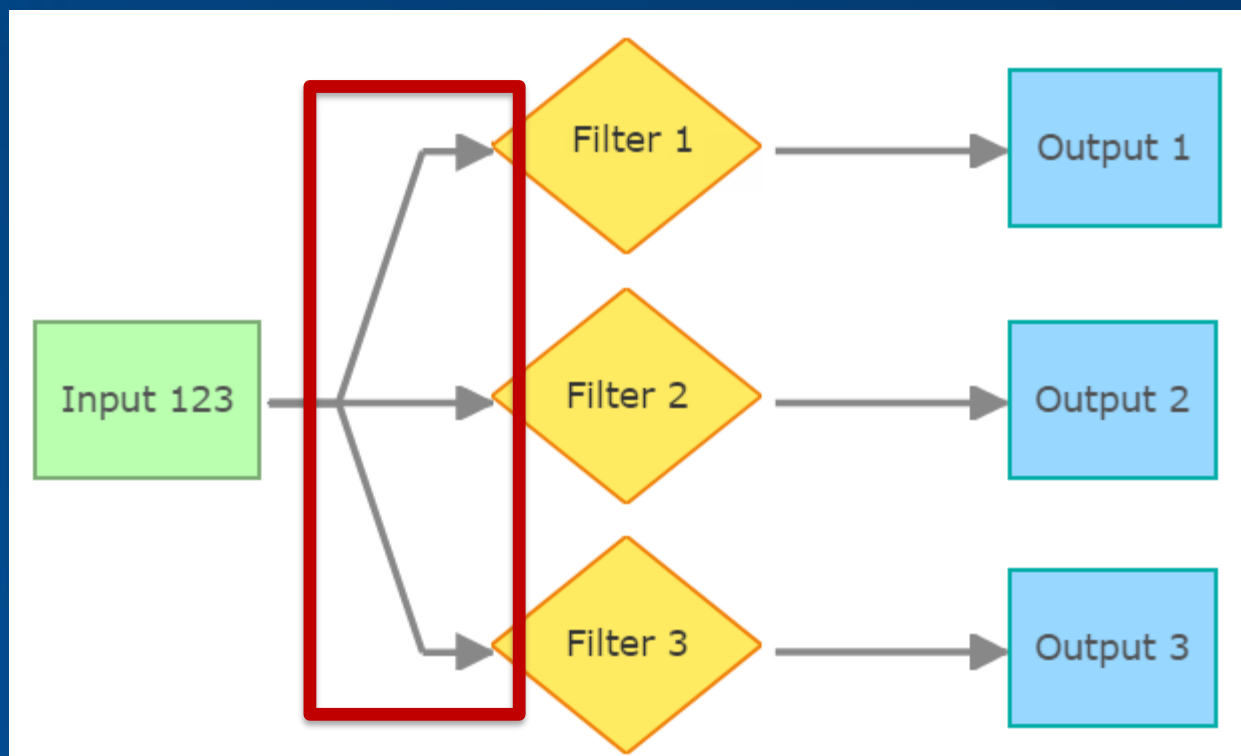
Service A



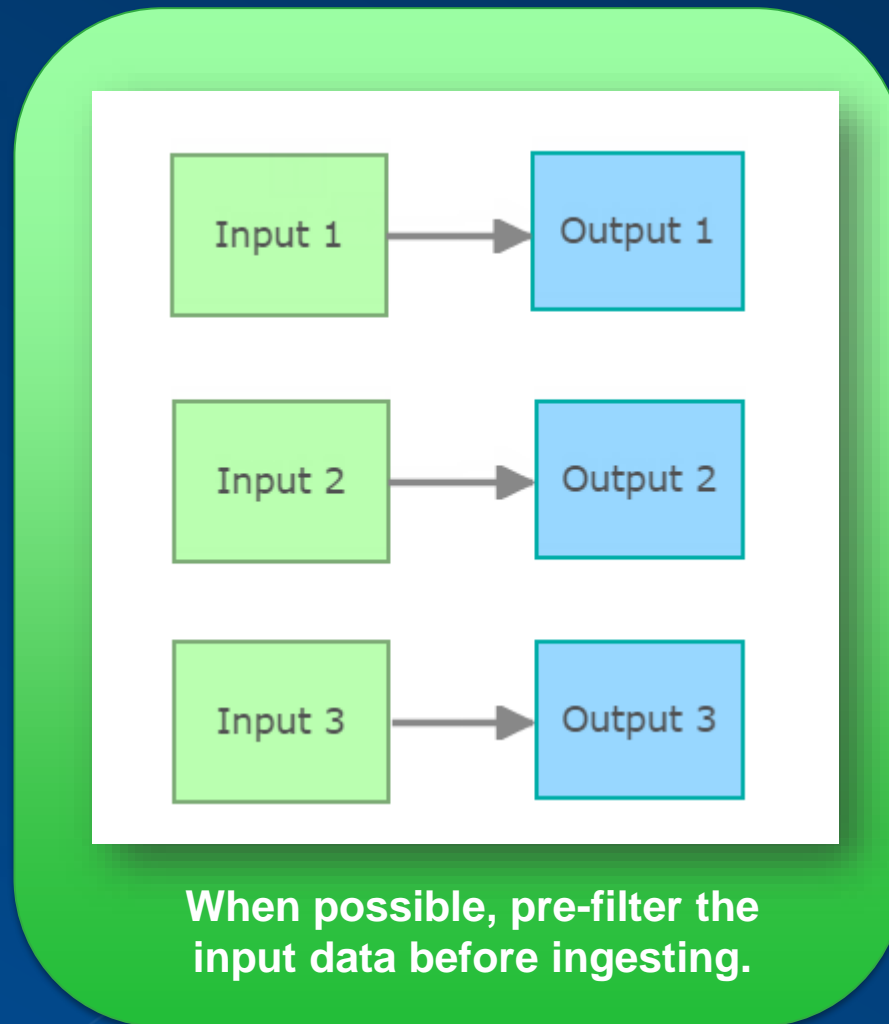
Service B

Service Design Considerations

which would you choose?



Each “branch” in a service contains the same event data. In this example, with three branches, it is creating 3X the volume of data.




When possible, pre-filter the input data before ingesting.

Service Design Considerations

not all components are created equally

A




**Which of these services
will process the fastest?
Slowest?**

B



C



Service Design Considerations

not all components are created equally


B



C



A



The first service only contains components that are utilizing the internal service cache, which allows for the fastest processing.

Service Design Considerations

not all components are created equally


B



C



A



The second service modifies the incoming event geometry which can be “costly”.

These types of requests are typically very quick but can be impacted by geometry complexity.


Service Design Considerations

not all components are created equally


B



C



A



The third service utilizes **Network Analyst** to return a “drive time” polygon which can significantly impact throughput.

Service Design Considerations

other recommendations

- **Configure Filters and/or Field Reducer Processors as early as possible in a service**
 - This reduces the volume / data size of the events being processed
 - Potentially simplifies service configuration “down stream”
- **Avoid Managed GeoEvent Definitions when possible**
 - These are “system owned” definitions whose lifecycle is entirely controlled by the processors
 - Editing or Deleting a processor will remove these definitions
 - If necessary copy generated definition and edit processor to look for it
- **Utilize the combination of Imported Definitions and Field Mapper Processor for Feature Service Outputs**
 - This ensures that all of the event data is being written in the correct format
 - Can also be used to update only a portion of the fields



Upgrade Planning

Upgrade Planning

what should be consider

- In-place Upgrade vs Clean Installation
 - When possible do a clean install
 - GeoEvent Server install and uninstalls very quickly
- Export Configuration & Global Settings from within GeoEvent Manager
 - Use time to remove any unused definitions or components
- Backup any configuration files that were modified in “/etc” folder
- Copy contents of “/deploy” folder (custom components)
- Delete contents of old site configuration (e.g. C:\arccgisserver\local\zookeeper)
- Install new version and import configurations



Troubleshooting

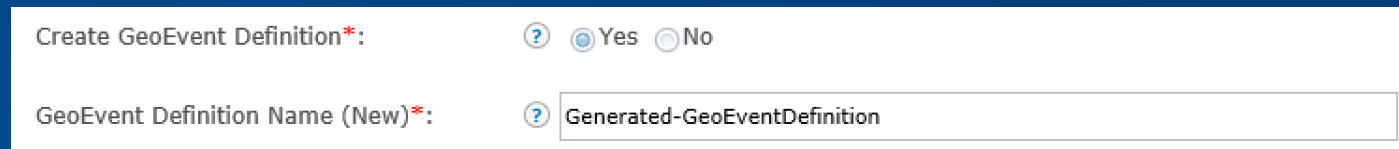
Troubleshooting

inputs

“I can’t get my data to come in...”

- **Check the definition**

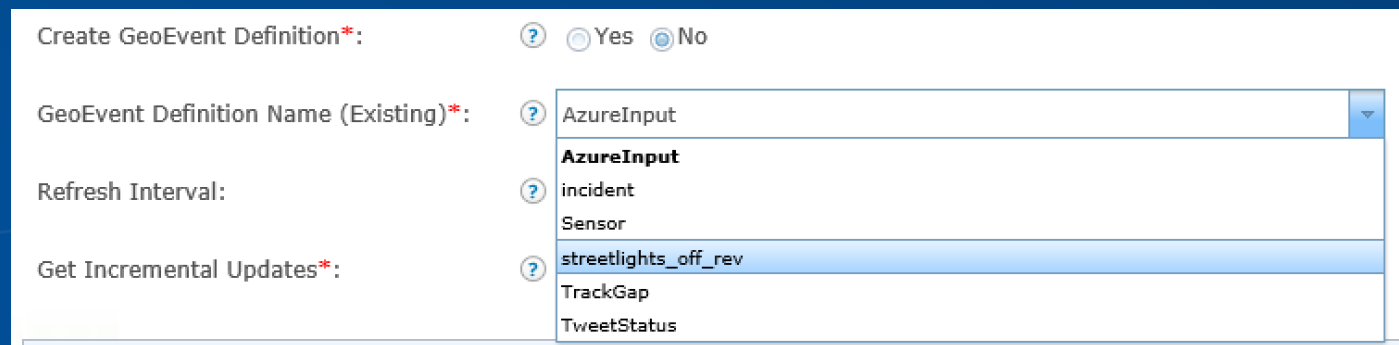
- Most input problems are with misconfigured schema (field names, data types, group structure)
- Try letting GeoEvent Server create definition for you...



Create GeoEvent Definition*: ☒ Yes ☐ No

GeoEvent Definition Name (New)*:

- ...but you will likely need to edit the definition and edit the input to use the one you modified



Create GeoEvent Definition*: ☒ Yes ☐ No

GeoEvent Definition Name (Existing)*:

Refresh Interval:

Get Incremental Updates*:

AzureInput

incident

Sensor

streetlights_off_rev

TrackGap

TweetStatus

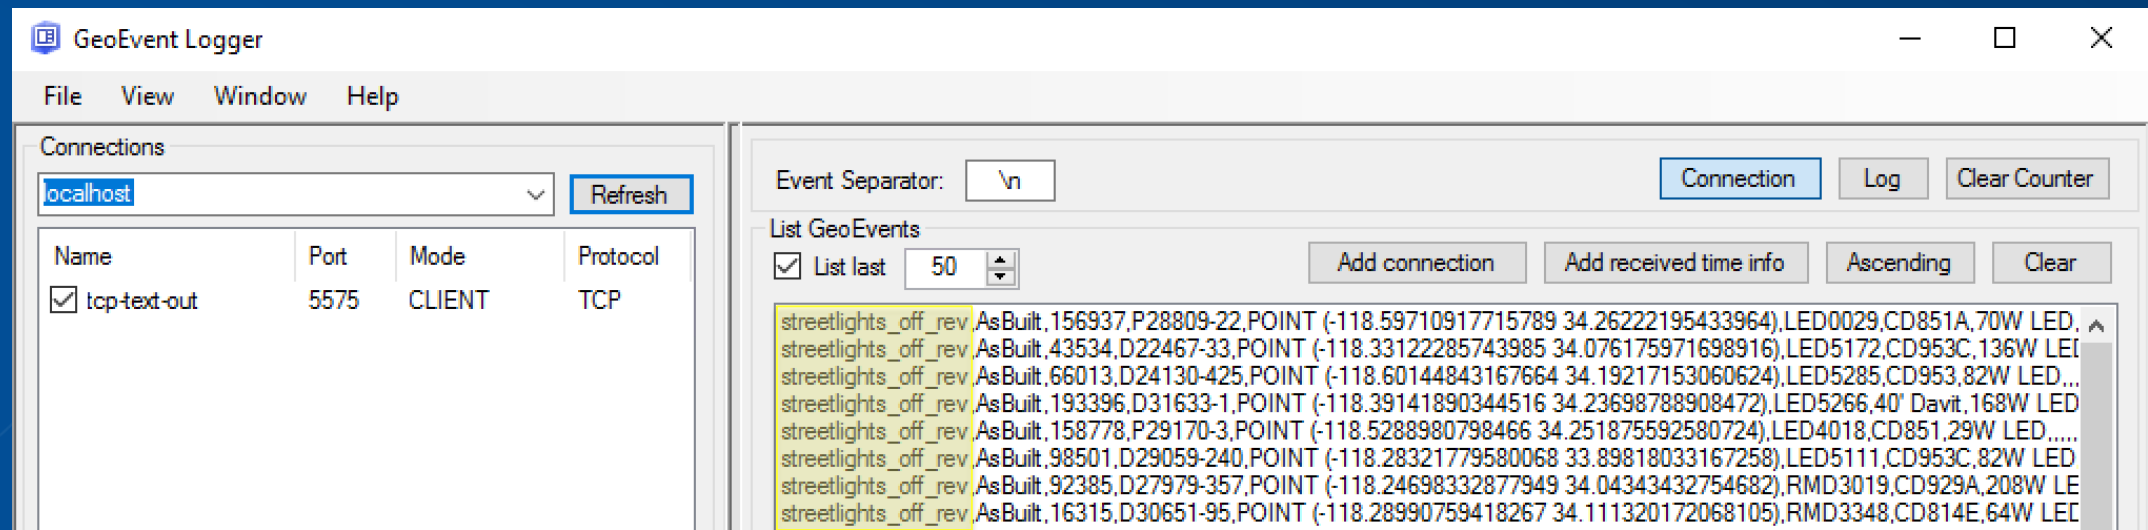
Troubleshooting

outputs

"I can't write my data to..."

- **Check the definition**

- Most output problems are with misconfigured schema (field names, data types, group structure)
- If possible import the definition from the service you are trying to write to
 - If an Esri Feature Service remove the reserved field names (e.g. ObjectID / OID)
- Use the GeoEvent Logger application to verify the expected output data



Troubleshooting

logging

"It doesn't work when I..."

- **Check the logs for errors**

- For critical problems the cause of the problem will likely display in the log message
- You may need to get more detail from a specific logger to isolate a cause

Logger	Message	Time	Level
com.esri.ges.messaging.jms.MessagingImpl	[GeoEvent-messaging] Consumer thread was interrupted.	Mar 1, 2018, 4:52:23 PM	ERROR

View Log

View Log Messages

Log Filter: ERROR

Logger

com.esri.ges.messaging.jms.MessagingImpl

com.esri.ges.messaging.jms.MessagingImpl

com.esri.ges.messaging.jms.MessagingImpl

Log Settings

Specify the level and logger for creating log messages.

Log level: DEBUG

Logger: com.esri.ges.messaging.jms.MessagingImpl

! The debug log level should only be used temporarily while debugging or troubleshooting as it may cause a significant decrease in performance.

Settings

Delete Logs

Time

Level

2018, 4:52:23 PM

ERROR

2018, 4:52:23 PM

ERROR

2018, 4:52:23 PM







ERROR

Troubleshooting

backup

“Everything was working yesterday but...” or “Someone accidentally deleted...”

- **Did you make a backup of your configuration?**
 - With ArcGIS GeoEvent Server 10.5 or newer, we did for you...

Name	Value	Action				
Automatic Backup Settings						
Backup Options	<table><tr><th>Back-up type</th><th>Properties</th></tr><tr><td>Daily</td><td>Time: 00:00</td></tr></table>	Back-up type	Properties	Daily	Time: 00:00	 
Back-up type	Properties					
Daily	Time: 00:00					
Enable Automatic Backups	true	 				
Max Backups	10	 				

Edit Folder

Register a folder.

Name: ?

Automatic Backups

X

Path: ?

C:\ProgramData\Esrri\GeoEvent

Real-Time & Big Data **Sessions!**

- **Real-Time GIS:** Best Practices
- **ArcGIS GeoEvent Server:** Configuring Real-Time Web Apps
- **ArcGIS GeoEvent Server:** An Introduction
- **Real-Time & Big Data GIS:** The Road Ahead
- **ArcGIS GeoEvent Server:** Applying Real-Time Analytics
- **Real-Time and Big Data:**
Leveraging the Spatiotemporal Big Data Store

Tue, 1:45 - 2:45 pm, 152 B

Wed, 2:45 - 3:45 pm, 152 B

Tue, 3:00 - 4:00 pm, 152 B

Wed, 5:15 - 6:15 pm, 152 B

Wed, 8:15 - 9:15 am, 152 B

Wed, 1:30 - 2:30 pm, 152 B

Wed, 4:00 – 5:00 pm, 152 B

Tue, 4:15pm - 5:15 pm, 152 B

Wed, 11:00 am -12:00 pm, 152 B

Print Your Certificate of Attendance

Print stations located in the 140 Concourse

Tuesday

12:30 pm – 6:30 pm

GIS Solutions Expo

Hall B

5:00 pm – 6:30 pm

GIS Solutions Expo Social

Hall B

Wednesday

10:30 am – 5:15 pm

GIS Solutions Expo

Hall B

6:30 pm – 9:00 pm

Networking Reception

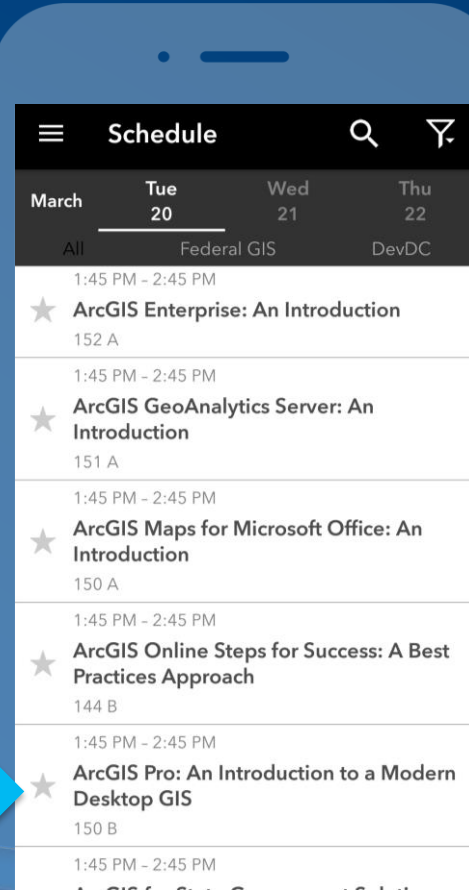
Smithsonian National Portrait Gallery

Please Take Our Survey in the Esri Events App

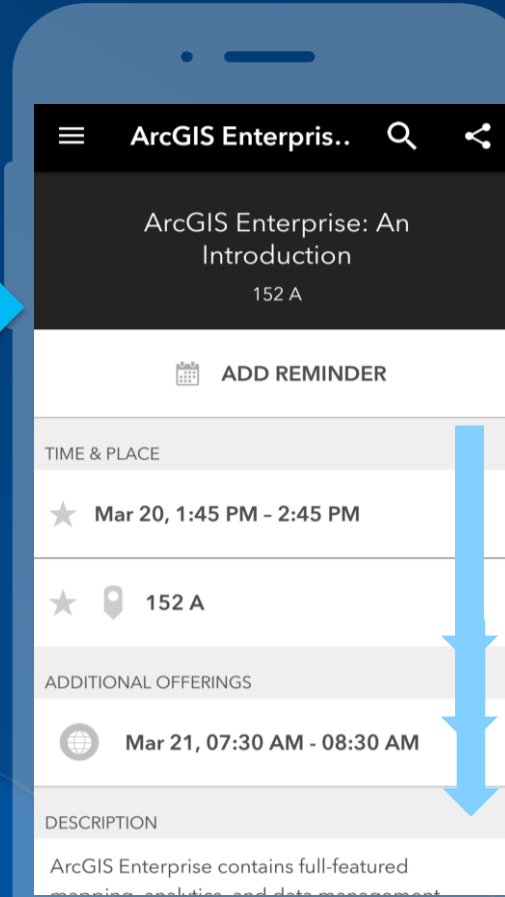
Download the Esri Events app and find your event



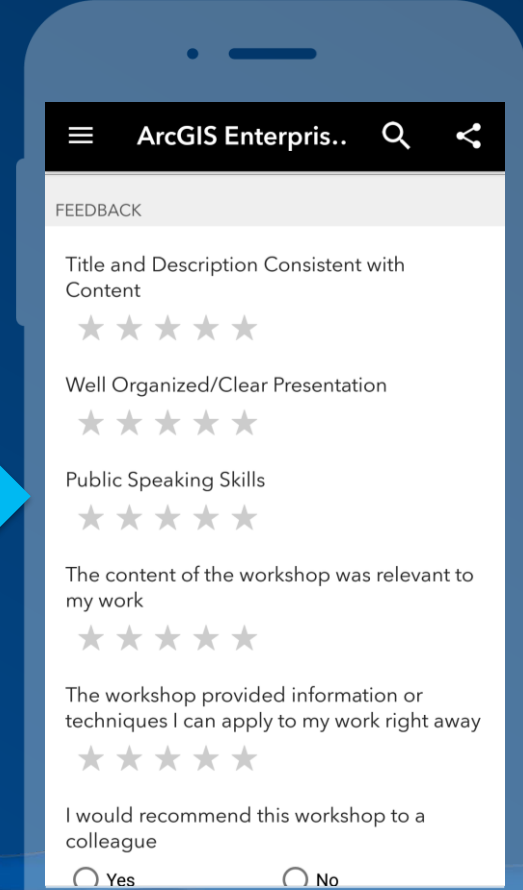
Select the session you attended



Scroll down to find the feedback section



Complete answers and select "Submit"





esri

THE
SCIENCE
OF
WHERE