

Development of IBM Watson with UIMA DUCC

Eddie Epstein
eae@apache.org

Apache UIMA
PMC Member and Committer


ApacheCon NA 2015





Presentation Outline


- What is Ducc
- Overview of the IBM-Jeopardy! Question-Answering System
- Interesting development problems
 - Solutions embodied in Ducc
- Fast cruise through Ducc's web interface





What is DUCC

- A Linux-based cluster controller designed specifically for UIMA
- Scales out any UIMA pipeline:
 - for high throughput, or
 - for low latency
- Uses CGroups to partition user processes
- Flexible Resource Management
- Extensive Web, CLI and API interfaces





What DUCC Does

- **Collection Processing Jobs**
 - Scale out a UIMA pipeline into multiple threads and processes, distribute collection as work items
- **Shared Services**
 - Manage life cycle of services, supporting dependencies with Jobs or other Services
- **Arbitrary Processes**
 - Launch arbitrary singleton processes or just provide a container to work

Motivations for DUCC

- Support Ongoing Watson Development
 - Take advantage of game playing hardware
 - Expanding development team
- Bring Functionality to Apache UIMA Community
 - Separate implementation from Watson code
 - Improve quality by targeting wide audience



Example Jeopardy Question




IN 1698, THIS COMET DISCOVERER TOOK A SHIP CALLED THE PARAMOUR PINK ON THE FIRST PURELY SCIENTIFIC SEA VOYAGE

Question Analysis

Keywords: 1698, comet, paramour, pink, ...
 AnswerType(comet discoverer)
 Date(1698)
 Took(discoverer, ship)
 Called(ship, Paramour Pink)
 ...


Primary Search



Candidate Answer Generation

- Isaac Newton
- Wilhelm Tempel
- HMS Paramour
- Christiaan Huygens
- Halley's Comet
- Edmond Halley
- Pink Panther
- Peter Sellers

Evidence Retrieval



Evidence Scoring

Lexical	Taxonomic	Spatial	...	Temporal
---------	-----------	---------	-----	----------

[0.58	0	-1.3	...	0.97]
[0.71	1	13.4	...	0.72]
[0.12	0	2.0	...	0.40]
[0.84	1	10.6	...	0.21]
[0.33	0	6.3	...	0.83]
[0.21	1	11.1	...	0.92]
[0.91	0	-8.2	...	0.61]
[0.91	0	-1.7	...	0.60]

Merging & Ranking

1. Edmond Halley (0.85)
2. Christiaan Huygens (0.20)
3. Peter Sellers (0.05)

...

Open Source Software Critical for Watson




Runtime

- Apache UIMA
- Indri Text Search (www.lemurproject.org/indri/)
- Apache Lucene (Text Search)
- Sesame (<http://aduna-software.com/technology/sesame>)
- Apache ActiveMQ (used by UIMA-AS)

During Development

- Eclipse (<https://eclipse.org>)
- Weka (<http://sourceforge.net/projects/weka/>)
- Apache Hadoop

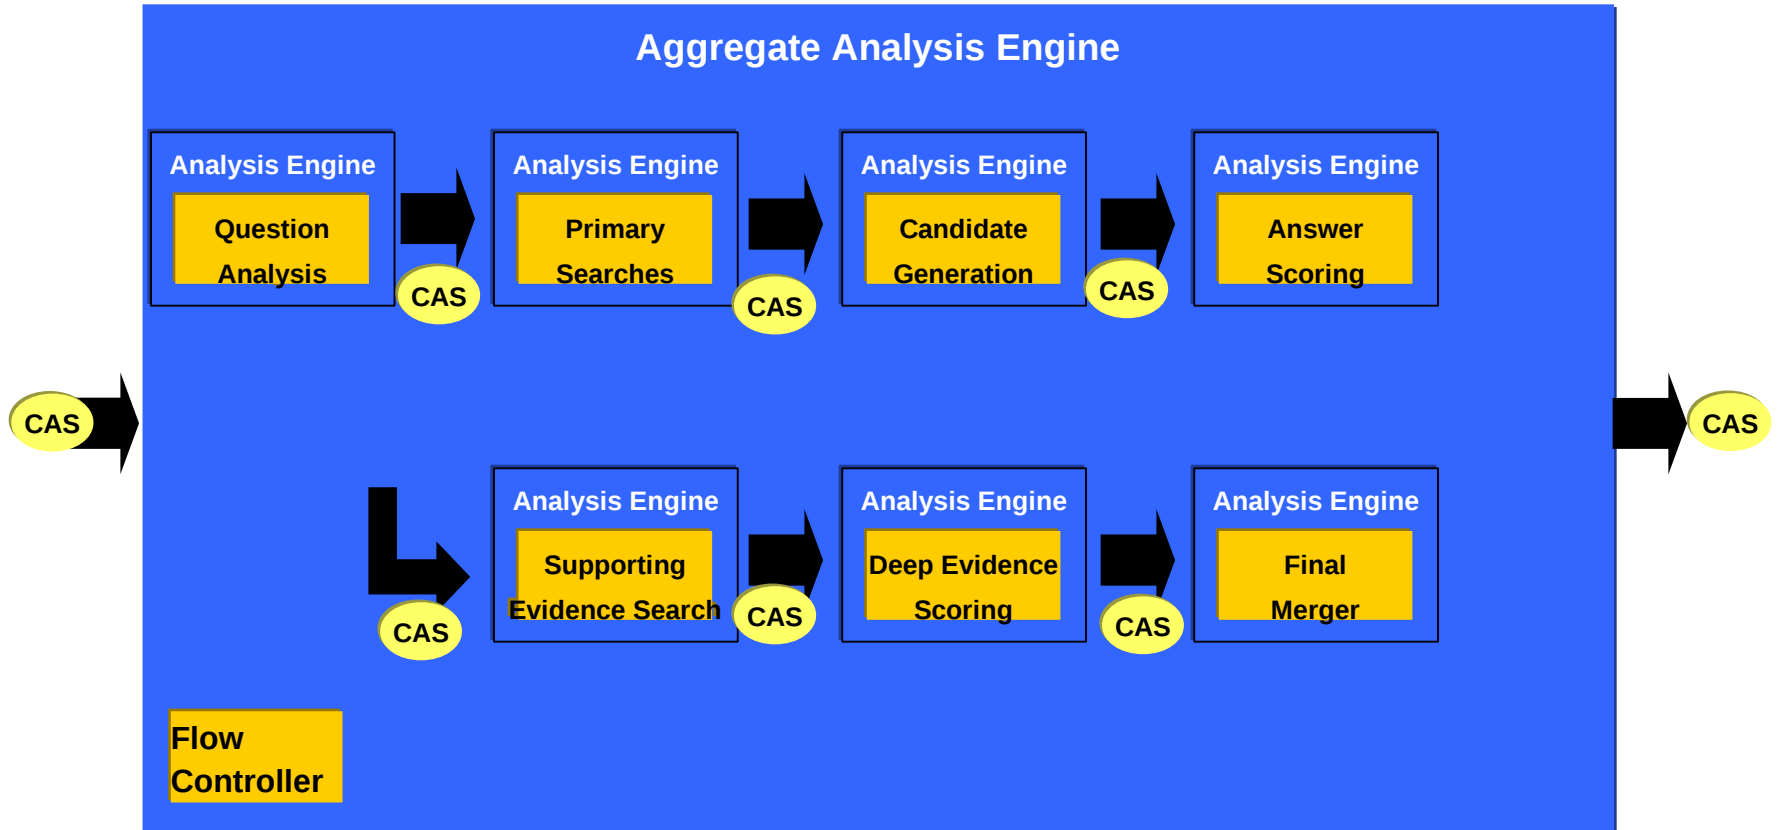
Watson's Knowledge for Jeopardy!



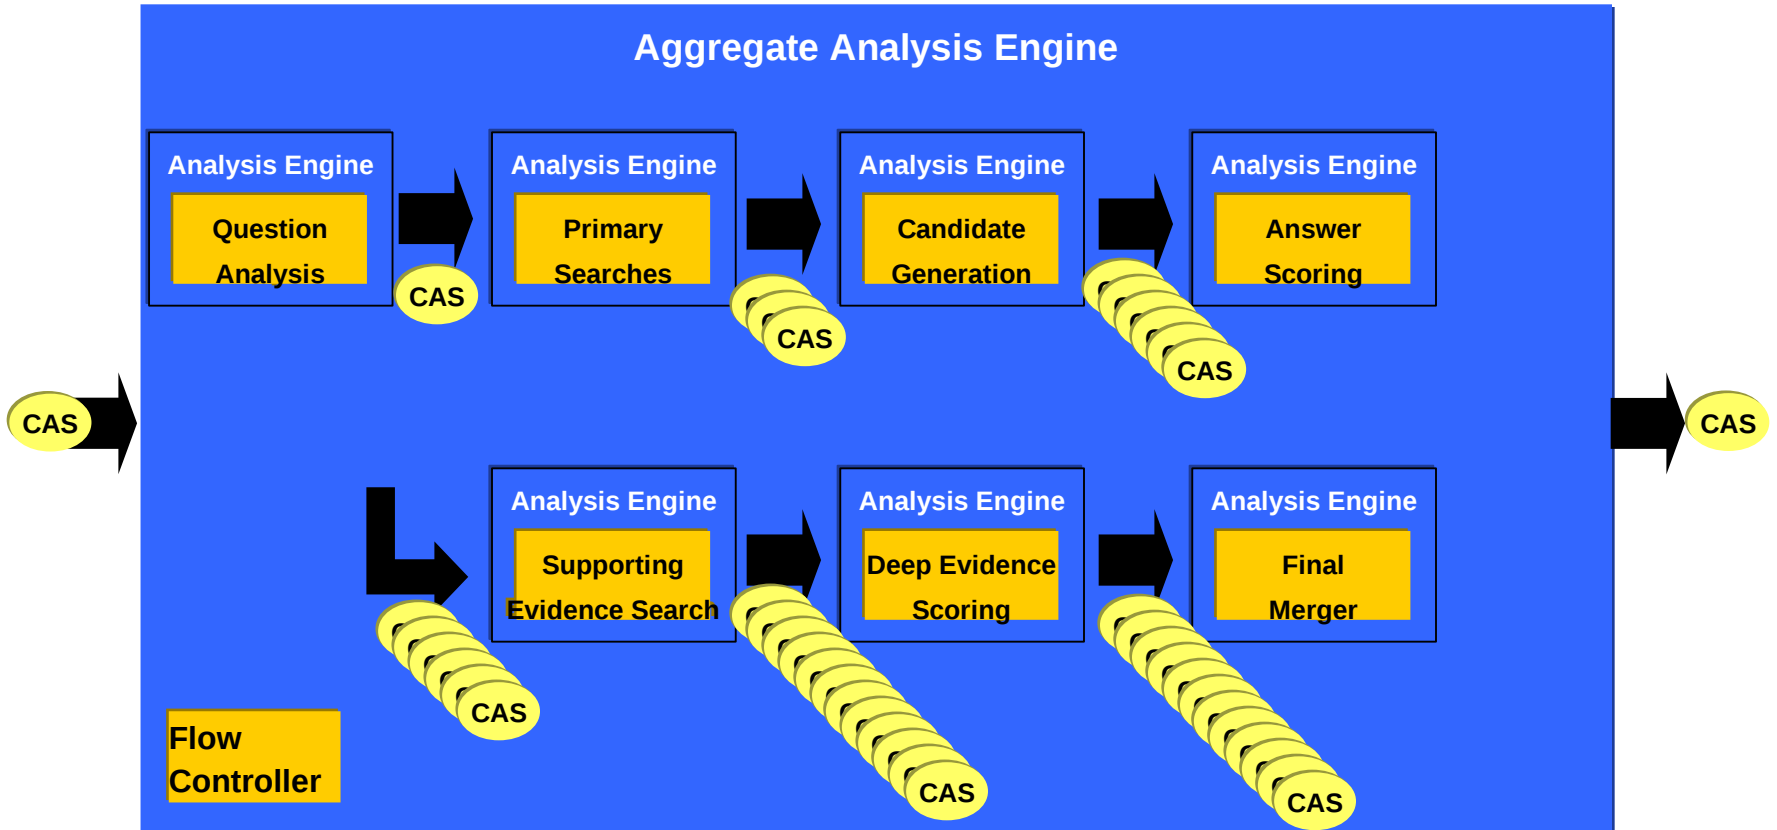
Watson has analyzed and stored the equivalent of about 1 million books (e.g., encyclopedias, dictionaries, news articles, reference texts, plays, etc)

Watson also uses structured sources such as WordNet and DBpedia

Watson on UIMA



Watson on UIMA – Data Flow





Problem – One Experiment

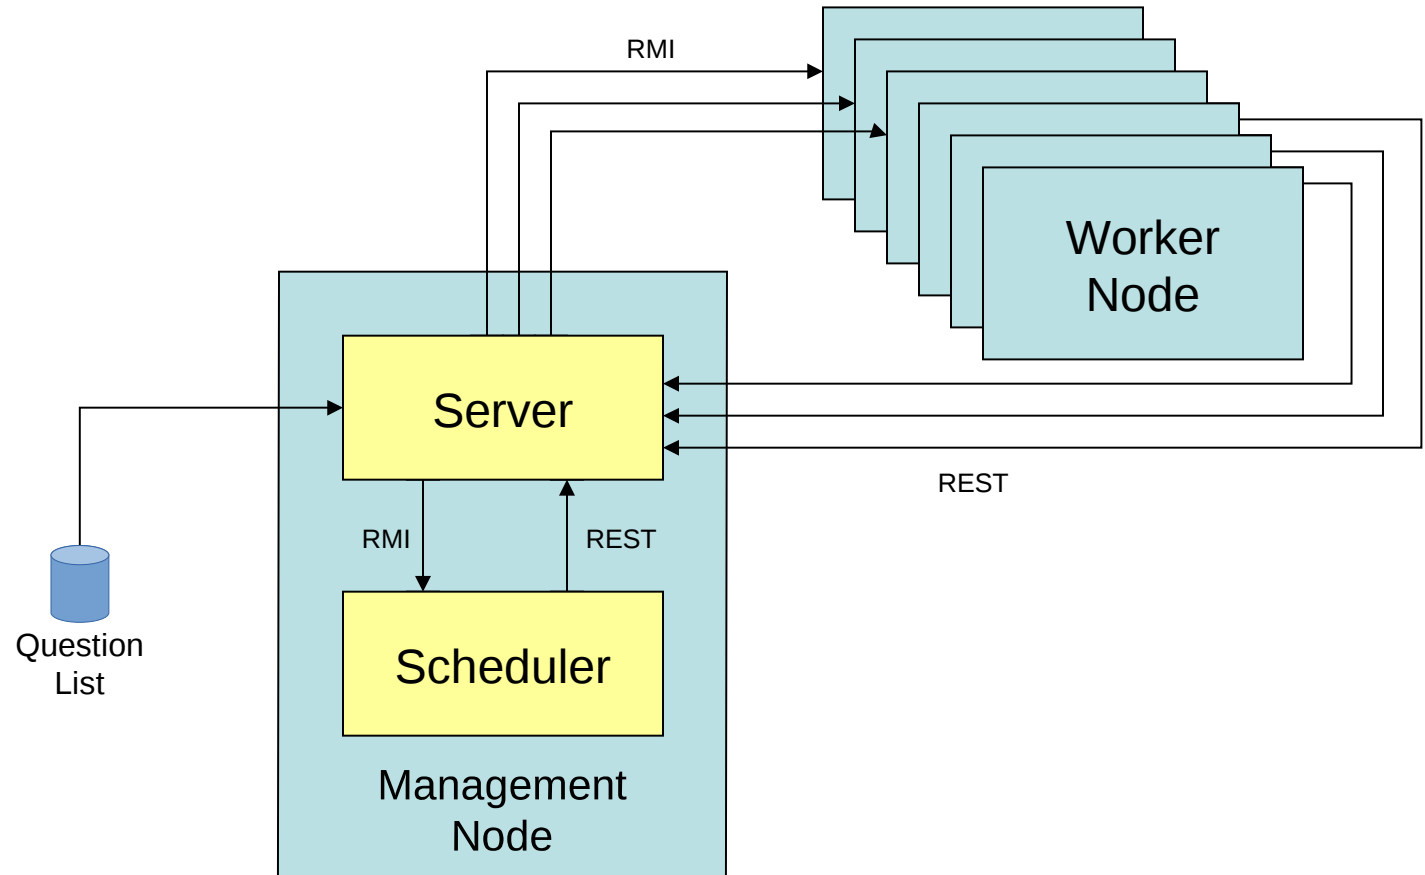
- **Average 2 hours per question**
 - Wide range of times
- **28GB Java Heap on 32GB Machines**
 - Large knowledge bases (e.g. Sesame in-memory store)
- **~1000 questions each**
 - To get statistically relevant results



Solution – One Experiment

- **Run parallel pipelines in multiple threads**
 - Share the large in-memory objects
 - Utilize the 8-cores in each machine
- **Replicate processes across machines**
 - Dynamically feed idle threads next question

BLADE Tool (before DUCC)

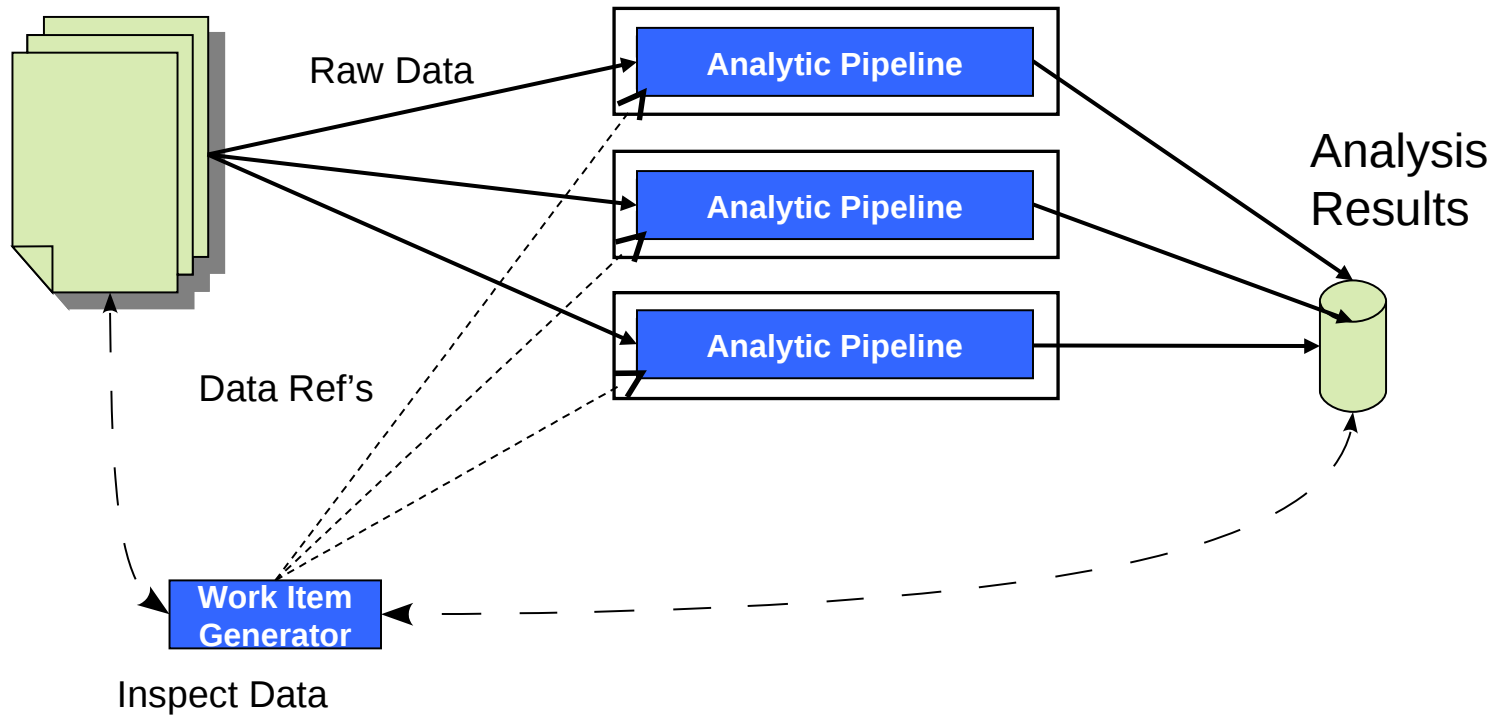


[http://domino.research.ibm.com/library/cyberdig.nsf/papers/152EF31994BD C3DC85257B1F005DE78F/\\$File/rc25356.pdf](http://domino.research.ibm.com/library/cyberdig.nsf/papers/152EF31994BD C3DC85257B1F005DE78F/$File/rc25356.pdf)



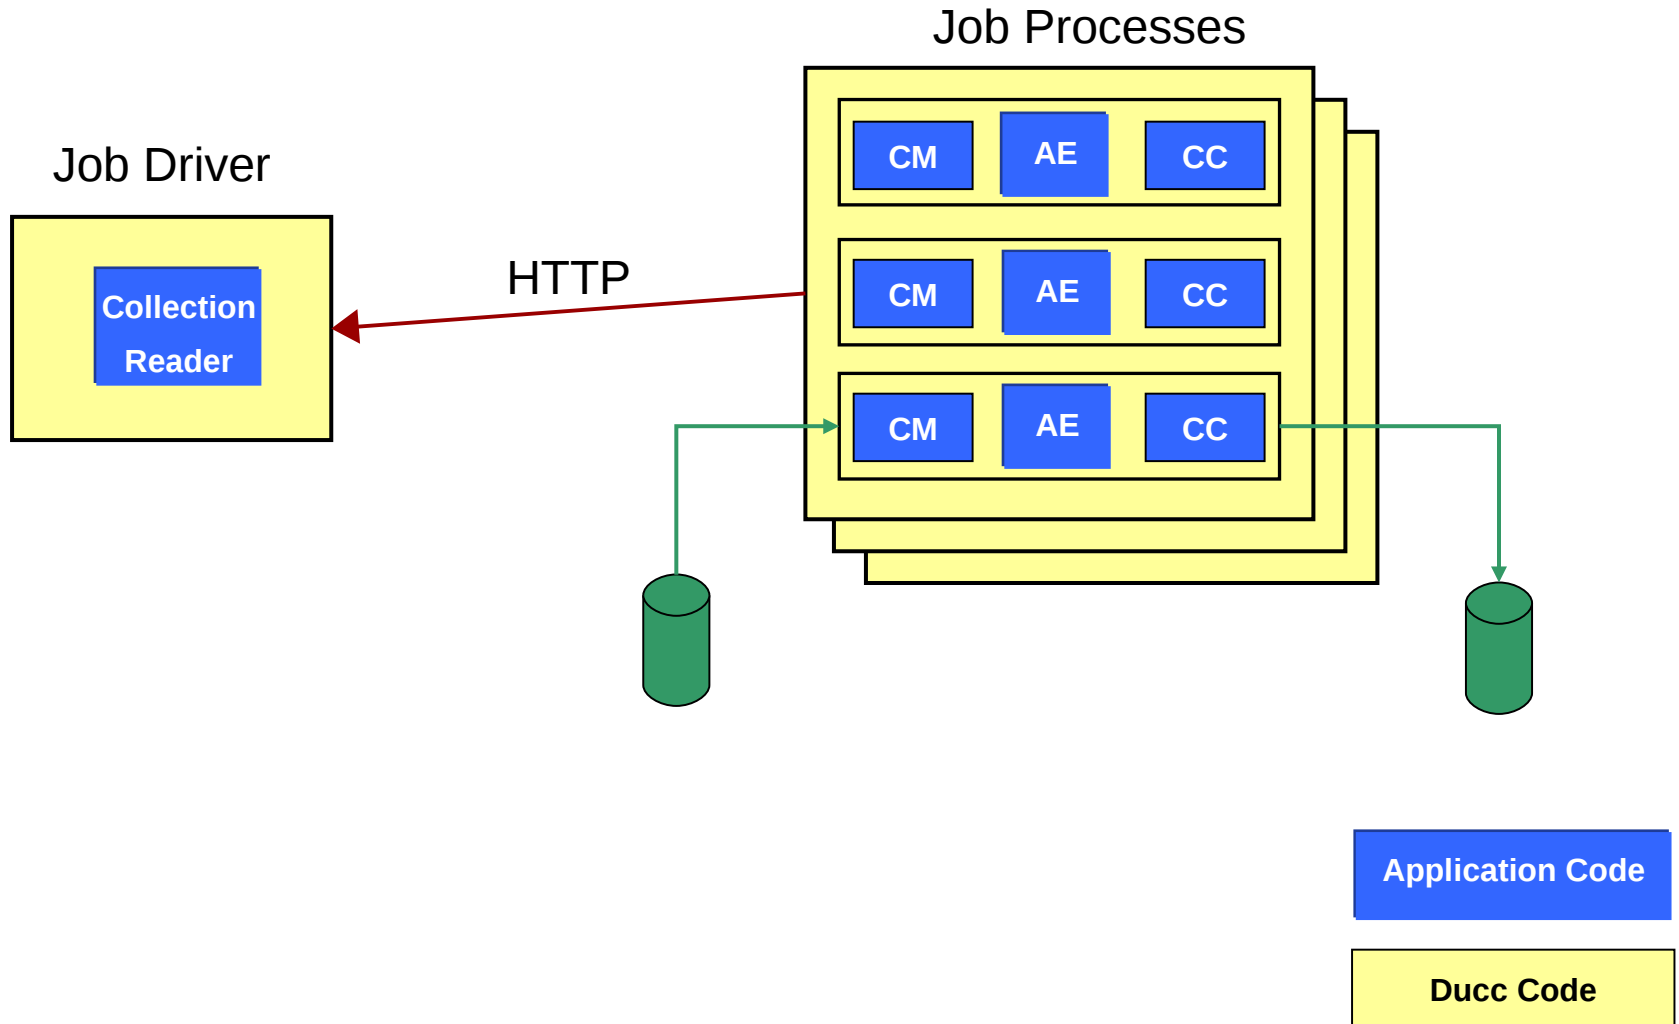
UIMA DUCC - Job Model

Collection of
Input Data



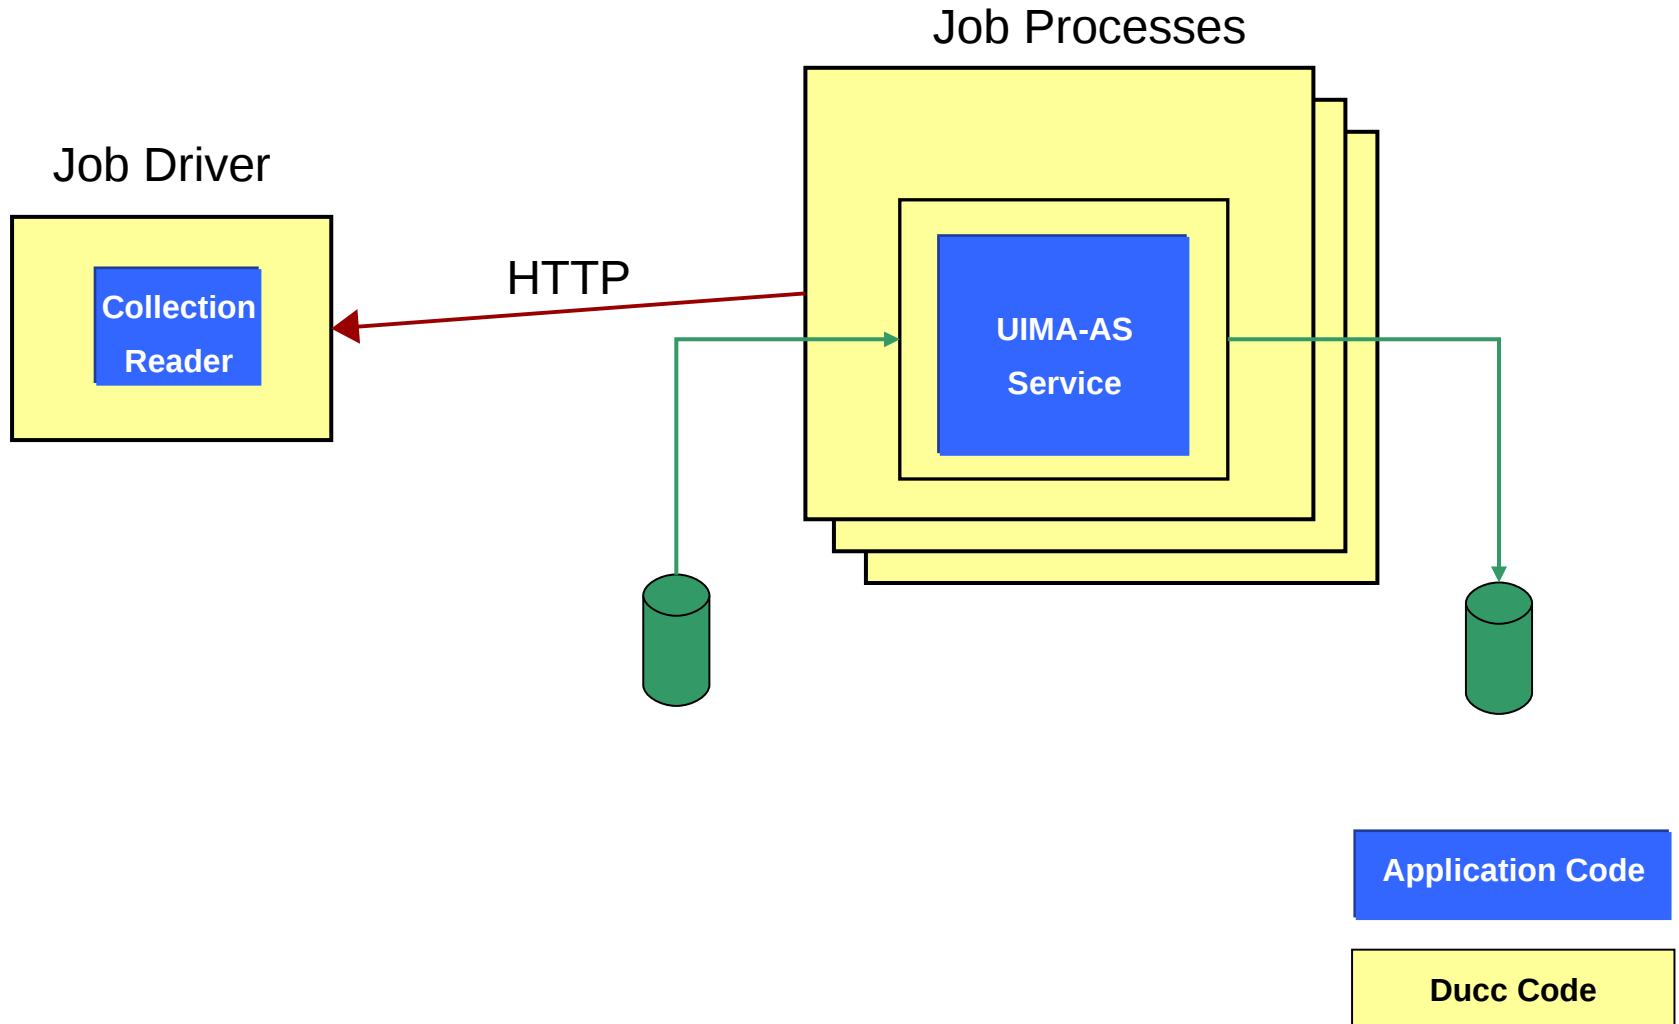


Job Model – Core UIMA Job



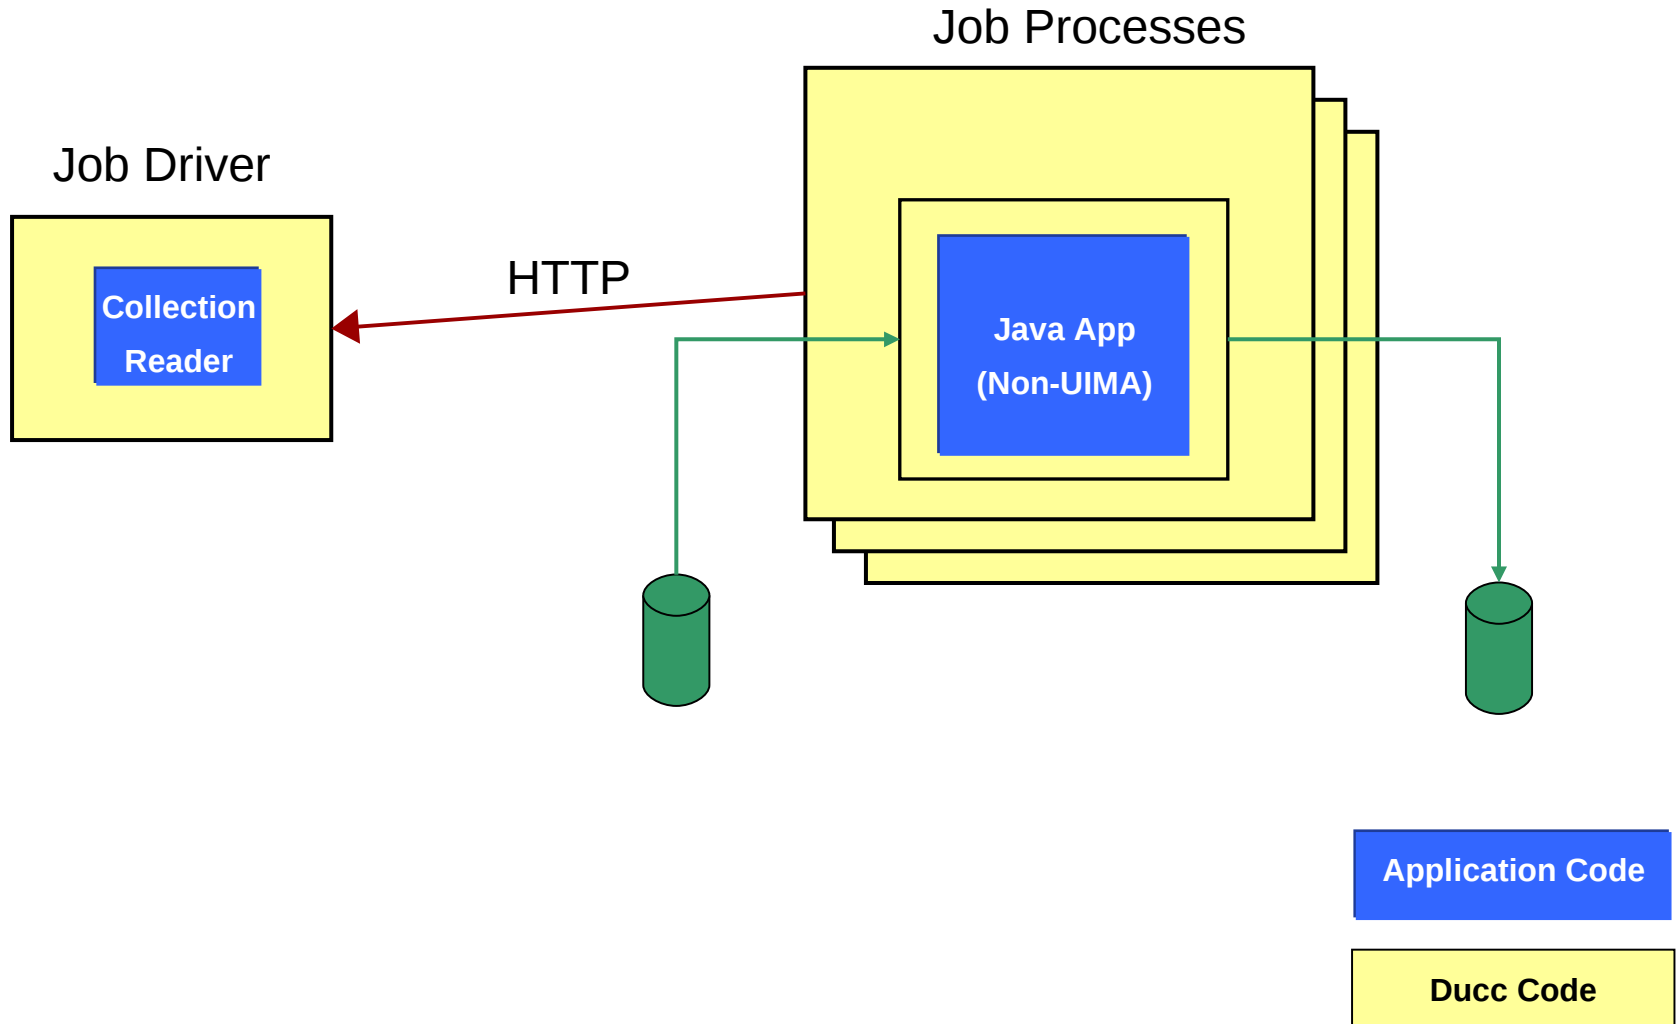


Job Model – UIMA-AS Job






Job Model – Custom Job





Job Debugging – all_in_one



All Job code deployed in a single thread in a single process for development & debug

Application Code

Ducc Code



Problem – 15 Researchers

- Personnel evaluated by their contribution to overall accuracy
 - With exceptions, e.g. reduce “stupid answers”
- Wanted their resource “fair share” NOW



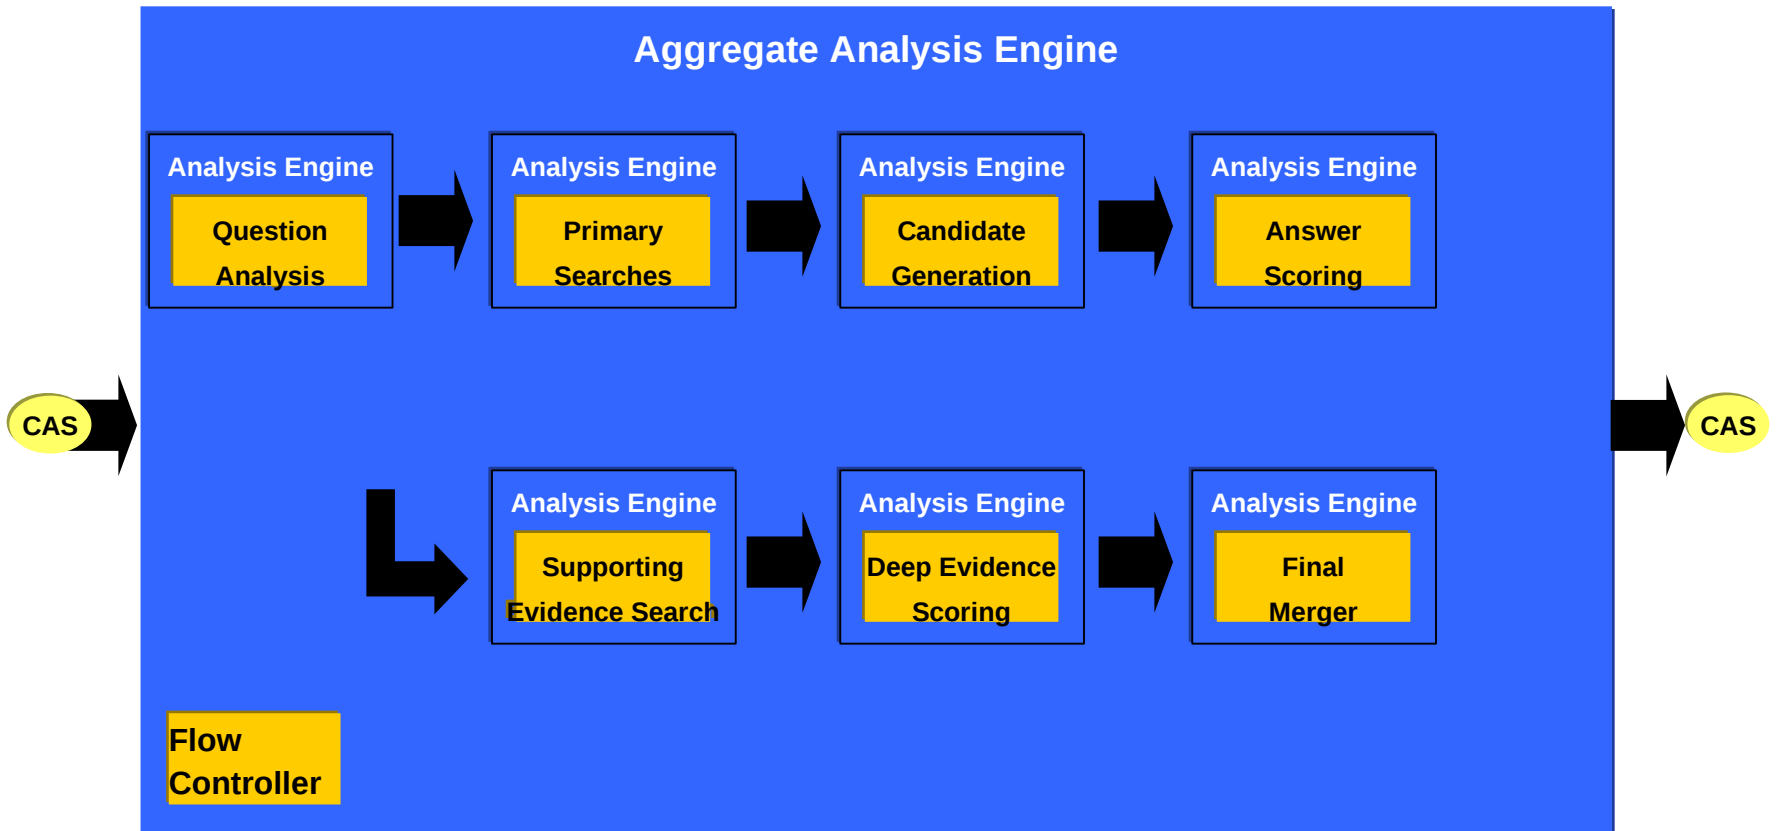


Solution – 15 Researchers

- **Preempt running processes**
 - Kill processes with least CPU investment
 - < 10% overhead for lost investment
- **Ramp up after successful initialization**
 - Saved more than preemption loses
- **Allow processes to be non-preemptable**
 - Reserve entire machines
 - Singleton processes (in CGroup containers)
 - Jobs

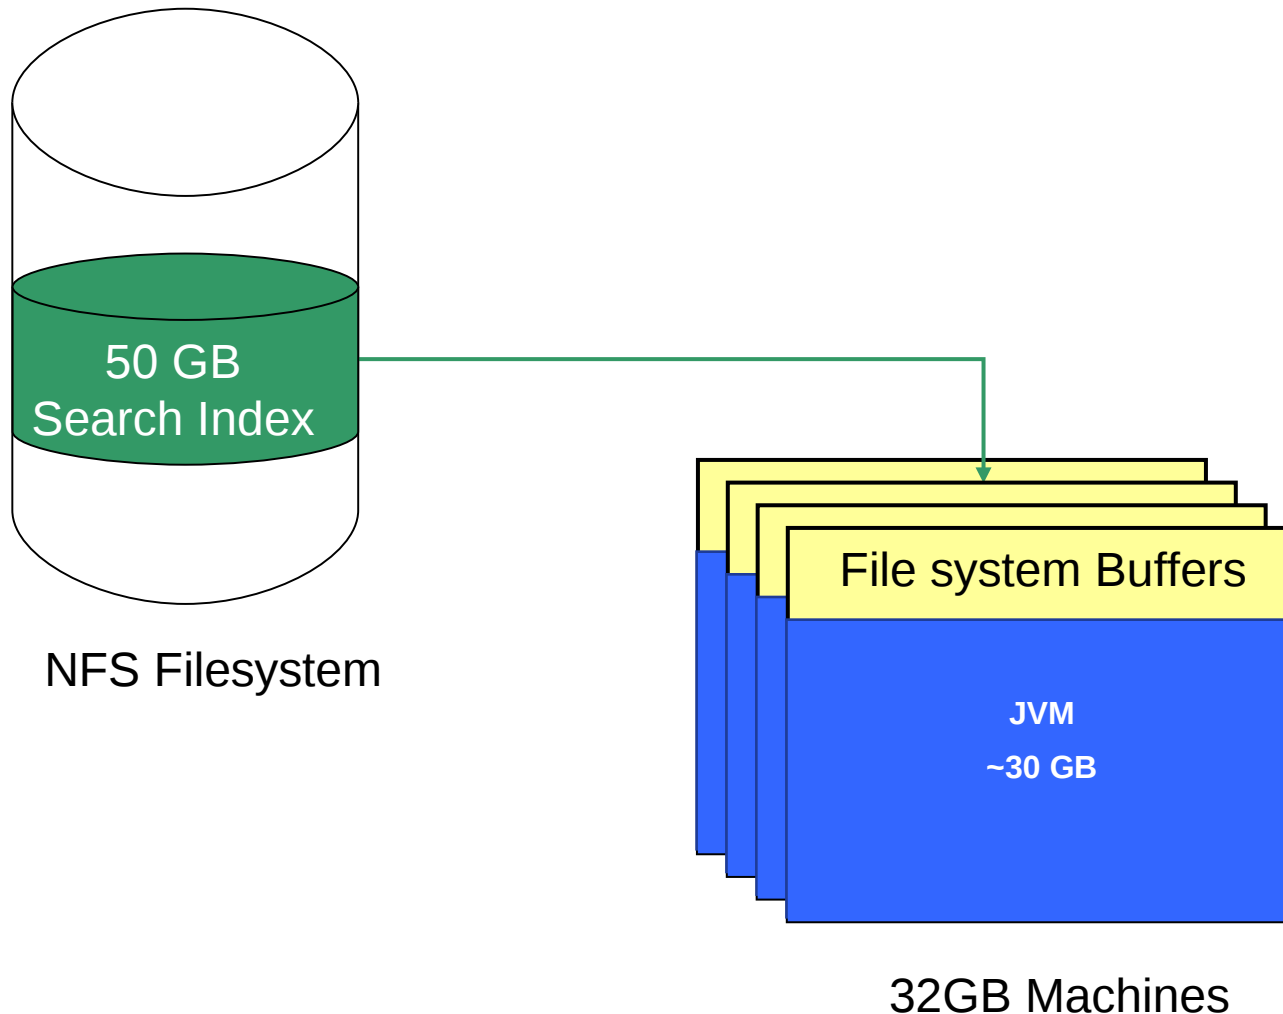


Watson on a 32GB Machine?

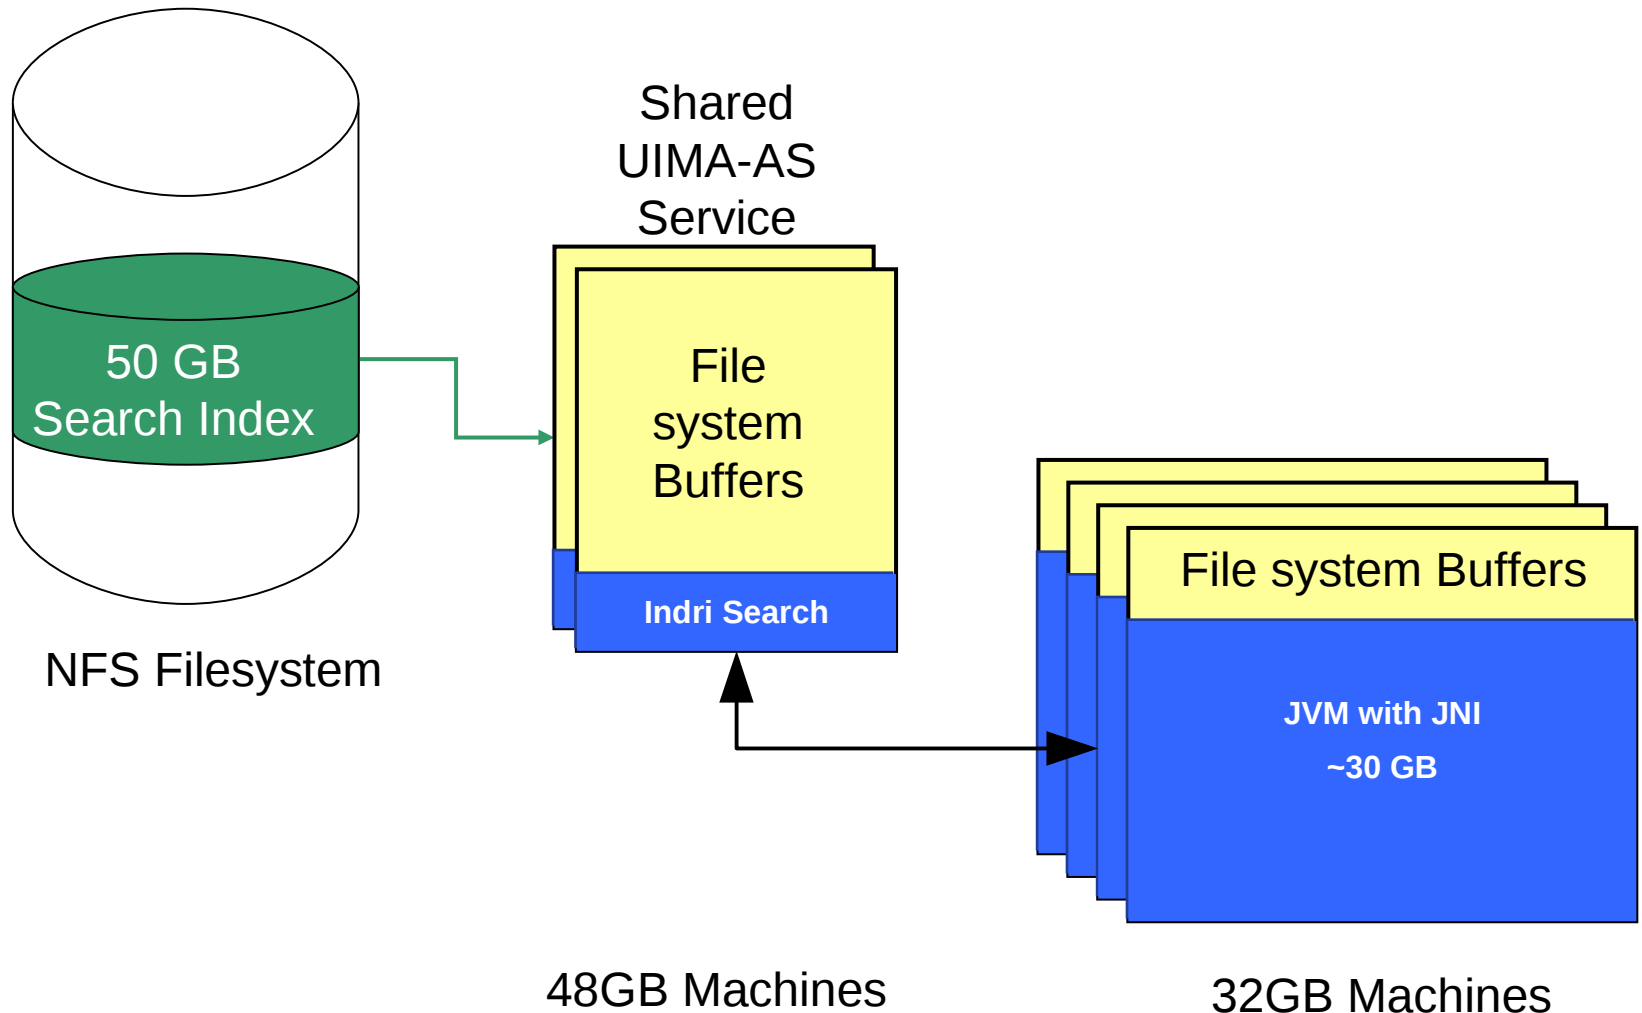


No, from the start some UIMA components were shared UIMA-AS services

Performance Bottleneck (Development Mode)



Services Improve Performance





Problem – Managing Services

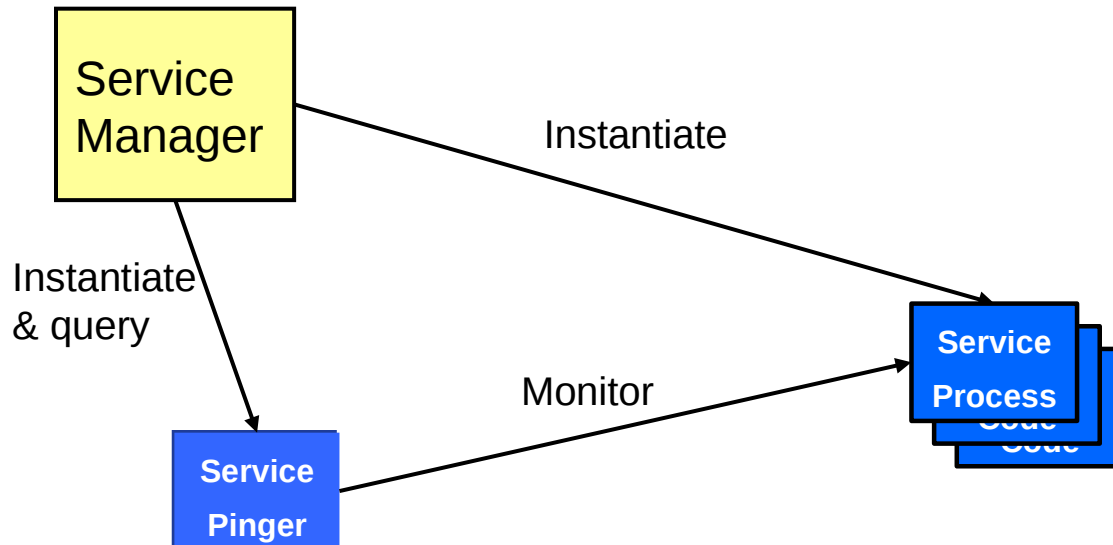
- Startup and number of instances manual
- Team had ~3 week sprints
 - Integrate changes and create new baseline
 - New indexes or code meant new services
 - Several baselines active concurrently



DUCC Services

- Service registry
- UIMA-AS or CUSTOM
 - Service “pinger” class required
 - Built-in pinger for UIMA-AS
- Always-on or start-on-demand
- Pinger interface supports autonomous instance management

DUCC Service




Application Code

DUCC Code

DUCC – Node Visualization



DUCC Web Demo



Backup if no Demo




DUCC Viz Page



Jobs **Reservations** **Services** **System** **Viz**

T E S T T E S T




Refresh

- Manual
- Automatic

Apache UIMA-DUCC
Updated: 2015.04.06 17:36:42 Mon
Utilization: 90.0%

Viz




Login | **Logged-out**

du
Distributed UIMA Cl
Version
Copyrights: © 2012-2015 The Apache

Sort By *Size* Name

Shares of size 15GB: 20, Jobs: 14, Services: 2, Managed Reservations: 1, Reservations: 1, Unoccupied: 2
RAM Total: 300GB, For shares: 300GB, Jobs: 210GB, Services: 30GB, Managed Reservations: 15GB, Reservations: 15GB, Unoccupied: 30GB





DUCC Job Page


Jobs

Reservations

Services

System

Viz



Refresh

Manual

Automatic

Apache UIMA-DUCC


Updated: 2015.04.06 17:44:02 Mon

Utilization: 80.0%

Users:

Max Records:

Jobs



Login | **Logged-out** | [Preferences](#) | [DuccBook](#)

ducc-mon

Distributed **UIMA Cluster Computing Monitor**

Version: 2.0.0-beta

Copyrights: © 2012-2015 The Apache Software Foundation and © 2011-2012 IBM

Jobs List

click column heading to sort

	Id	Start	Duration	User	Class	State	Reason or extraordinary status	Services	Processes	Init Fails	Run Fails	Pgin	Swap	Size	Total	Done	Error	Dis- patch	Retry	Pre- empt	Description
<input type="button" value="Terminate"/>	176	2015.04.06 17:23:57 Mon	20:13 {+16:29}	jellyfish	normal	Running		0	2	0	0	0	0.0	30	1001	577	0	16	0	0	type #2 item count "medium"
<input type="button" value="Terminate"/>	175	2015.04.06 17:23:06 Mon	21:04 {+21:30}	albatross	normal	Running		0	4	0	0	0	0.0	30	1001	451	0	16	0	0	type #2 item count "medium"
	179	2015.04.06 17:35:15 Mon	07:16	hummingbird	low	Completed	EndOfJob	1	0	0	0	0	0.0	30	101	101	0	0	0	0	type #1 item count "small"
	177	2015.04.06 17:24:41 Mon	07:19	tapir	normal	Completed	EndOfJob	2	0	0	0	0	0.0	15	101	101	0	0	0	0	type #1 item count "small"



Job Details - Processes

- Jobs
- Reservations
- Services
- System
- Viz



Refresh

Manual

Automatic

Apache UIMA-DUCC
 Updated: 2015.04.06 17:46:12 Mon
 Utilization: 80.0%

[Job Details](#)



Login | **Logged-out** | [Preferences](#) | [DuccBook](#)

ducc-mon

Distributed UIMA Cluster Computing Monitor

Version: 2.0.0-beta

Copyrights: © 2012-2015 The Apache Software Foundation and © 2011-2012 IBM

Id: 176 State: Running Workitems: 1001 Done: 643 Error: 0 Dispatch: 16

- Processes
- Work Items
- Performance
- Specification
- Files

Processes List

click column heading to sort


Id	Log	Size	Host Name	PID	State Scheduler	Reason Scheduler or extraordinary status	State Agent	Reason Agent	Exit	Time Init	Time Run	Time GC	PgIn	Swap	%CPU	RSS	Time Avg	Time Max	Time Min	Done	Error	Dis-patch	Retry	Pre-empt	JConsole URL
0	176-JD-uima-ducc-demo-3-12267.log	6.78	uima-ducc-demo-3	12267	Allocated		Running			00	22:13	00	0	0.0	0.0	0.1	25	60	10	643	0	16	0	0	service:jmx:rmi:///jndi/rmi://uima-ducc-vm:11018/jmxrmi
11	176-UIMA-uima-ducc-demo-10-13880.log	0.42	uima-ducc-demo-10	13880	Allocated		Running			25	20:48	01	0	0.0	0.0	0.1	35	60	10	272	0	8	0	0	service:jmx:rmi:///jndi/rmi://uima-ducc-vm:11024/jmxrmi
14	176-UIMA-uima-ducc-demo-9-17754.log	0.39	uima-ducc-demo-9	17754	Allocated		Running			20	18:53	01	0	0.0	0.0	0.1	34	60	10	252	0	8	0	0	service:jmx:rmi:///jndi/rmi://uima-ducc-vm:11027/jmxrmi
16	176-UIMA-uima-ducc-demo-10-20757.log	0.20	uima-ducc-demo-10	20757	Deallocated	Forced	Stopped	Deallocated	ExitCode=1	08	17:26	00	0	0.0	1.0	0.1	37	59	10	119	0	0	0	0	



Job Details - Performance

Jobs
Reservations
Services
System
Viz

T
E
S
T




T
E
S
T

Refresh
 Manual
 Automatic

Apache UIMA-DUCC
Updated: 2015.04.06 17:46:12 Mon
Utilization: 80.0%

Job Details



Id: 176 State: Running Workitems: 1001 Done: 643 Error: 0 Dispatch: 16


Name	Total ddd:hh:mm:ss	% of Total	Avg hh:mm:ss/workitem	Min hh:mm:ss/workitem	Max hh:mm:ss/workitem
Summary	06:17:43.5	100.0	35.2	05.1	01:00.2
DefaultAE	06:17:41.8	100.0	35.2	10.0	59.9
DUCC Job Flow Controller	01.7	0.0	00.0	00.0	00.2




DUCC Service Page

[Jobs](#)
[Reservations](#)
[Services](#)
[System](#)
[Viz](#)

TEST



TEST



Refresh

 Manual
 Automatic

Apache UIMA-DUCC


Updated: 2015.04.06 17:35:02 Mon

Utilization: 90.0%

Users:

Max Records:

Services



[Login](#) | [Logged-out](#) | [Preferences](#) | [DuccBook](#)

ducc-mon

Distributed UIMA Cluster Computing Monitor

Version: 2.0.0-beta

Copyrights: © 2012-2015 The Apache Software Foundation and © 2011-2012 IBM

Services Definitions List

click column heading to sort

	Id	Name	State	Last Use	Instances	Deployments	Start State	User	Class	Pgin	Swap	Size	Jobs	Ser-vices	Reser-vations	Description
<input type="button" value="Stop"/>	11	CUSTOM:localhost:7175	Available		1	1	Autostart	degenaro		0	0.0	0	0	0	0	Custom Service Type auto-start
<input type="button" value="Stop"/>	10	UIMA-AS:FixedSleepAE_4.tcp://localhost:61617	Available		1	1	Reference	zebra	fixed	1	0.0	15	0	0	0	Test Service 4 Type on-demand
	9	UIMA-AS:FixedSleepAE_3.tcp://localhost:61617	Stopped		1	0	Stopped	moose	fixed	0	0.0	15	0	0	0	Test Service 3 Type on-demand
	8	UIMA-AS:FixedSleepAE_2.tcp://localhost:61617	Stopped		1	0	Stopped	squirrel	fixed	0	0.0	15	0	0	0	Test Service 2 Type on-demand
	7	UIMA-AS:FixedSleepAE_1.tcp://localhost:61617	Stopped		1	0	Stopped	kiwi	fixed	0	0.0	15	0	0	0	Test Service 1 Type auto-start
	6	UIMA-AS:FixedSleepAE_0.tcp://localhost:61617	Stopped		1	0	Stopped	chinchilla	fixed	0	0.0	15	0	0	0	Test Service 0 Type on-demand

DUCC Reservation Page



[Jobs](#)
[Reservations](#)
[Services](#)
[System](#)
[Viz](#)

Refresh
 Manual
 Automatic

Apache UIMA-DUCC

Updated: 2015.04.06 17:49:32 Mon

Utilization: 90.0%

Users:

Max Records:

Reservations

[Login](#) | [Logged-out](#) | [Preferences](#) | [DuccBook](#)
ducc-mon
 Distributed UIMA Cluster Computing Monitor
 Version: 2.0.0-beta
Copyrights: © 2012-2015 The Apache Software Foundation and © 2011-2012 I

Reservations List

click column heading to sort

	Id	Start	Duration	User	Class	Type	State	Reason	Allocation	User Processes	Size	Host Names	Description
<input type="button" value="Terminate"/>	174	2015.04.06 17:20:30 Mon	29:09	quail	fixed	Managed	Running		1	1	15	uima-ducc-demo-1 #2	type "managed", end by user cancel
<input type="button" value="Terminate"/>	172	2015.04.06 17:17:55 Mon	31:44	System	JobDriver	Unmanaged	Assigned		1	0	1	uima-ducc-demo-3	Job Driver
	169	2015.04.06 16:41:57 Mon	26	bonobo	fixed	Managed	Completed	ProgramExit	1	0	15	uima-ducc-demo-4 #3	type "managed", end by program exit
	168	2015.04.06 16:00:41 Mon	22	eagle	fixed	Managed	Completed	ProgramExit	1	0	15	uima-ducc-demo-4 #3	type "managed", end by program exit
	165	2015.04.06 15:10:09 Mon	47:01	eagle	fixed	Managed	Completed	CanceledByUser	1	0	15	uima-ducc-demo-4 #2	type "managed", end by user cancel
	156	2015.04.06 14:23:06 Mon	27	urchin	fixed	Managed	Completed	ProgramExit	1	0	15	uima-ducc-demo-4 #3	type "managed", end by program exit
	152	2015.04.06 13:39:33 Mon	20	frog	fixed	Managed	Completed	ProgramExit	1	0	15	uima-ducc-demo-8 #3	type "managed", end by program exit
	144	2015.04.06 12:38:59 Mon	58:01	vicuna	fixed	Managed	Completed	CanceledByUser	1	0	15	uima-ducc-demo-1 #2	type "managed", end by user cancel
	132	2015.04.06 11:59:44 Mon	19	kiwi	fixed	Managed	Completed	ProgramExit	1	0	15	uima-ducc-demo-1 #3	type "managed", end by program exit

Thank You

