



# ACCELERATE AI DATA PIPELINE

With FlashStack by Cisco and Pure Storage




# COMPLEXITIES OF AI IN PRODUCTION



## INGEST

From sensors, machines,  
& user generated




COPY &  
TRANSFORM

## CLEAN & TRANSFORM

Label, anomaly detection,  
ETL, prep, stage



CPU Servers




COPY &  
TRANSFORM

## EXPLORE

Quickly iterate to  
converge on models



GPU Servers




COPY &  
TRANSFORM

## TRAIN


Run for hours to days in  
production cluster



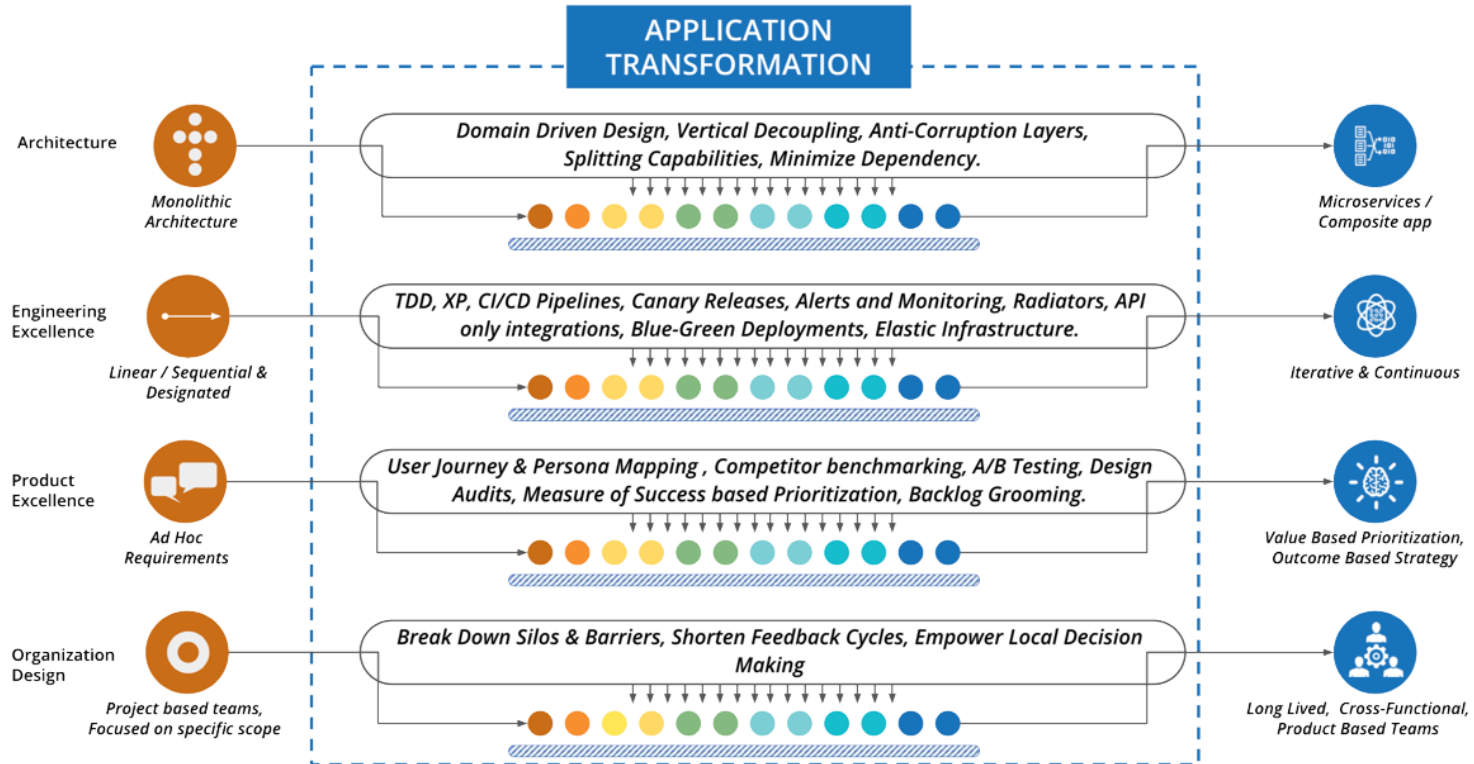
GPU Production Cluster



# WITHIN EACH AI STAGE



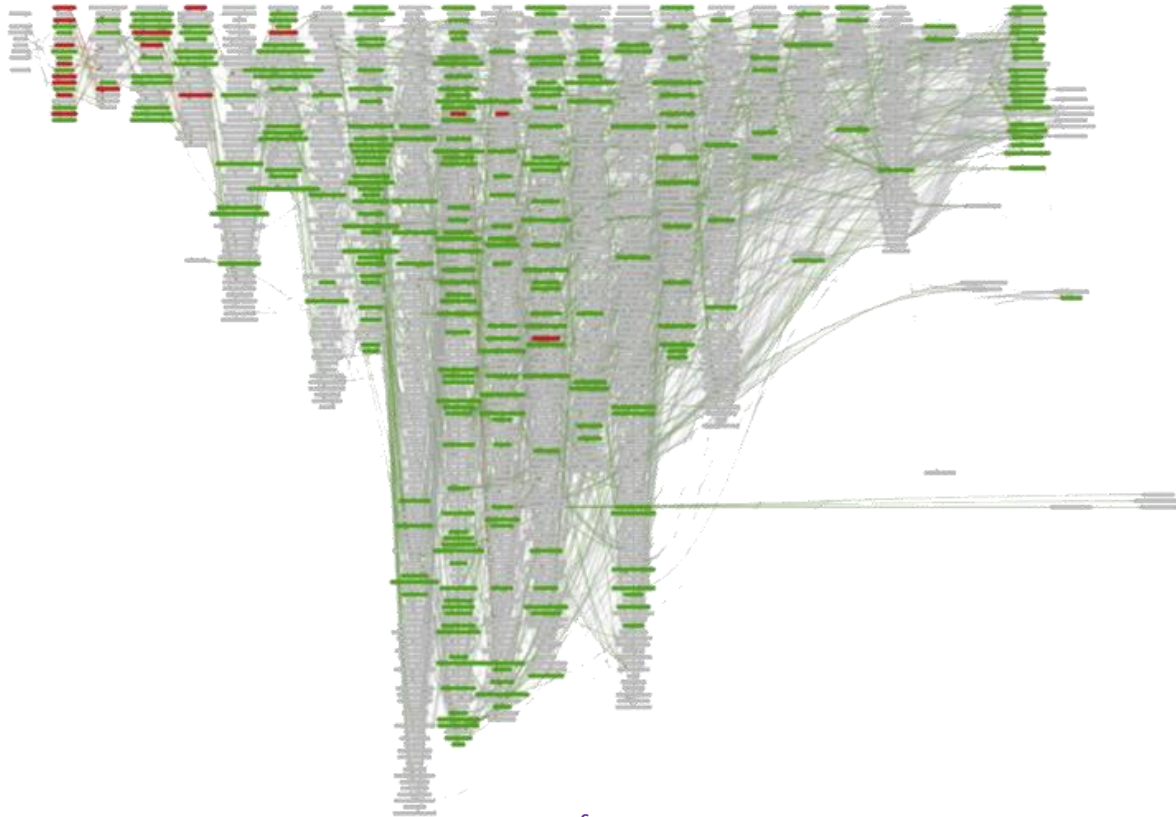
# DEPLOYED VIA MICROSERVICES



Source:  
ThoughtWorks – Mike Beaven, Digital Dialogue 2019



# COMPLEXITY AT SCALE



Source:  
[Medium – Microservices Visualization Tool](#)



# ANNOUNCING FLASHSTACK FOR AI

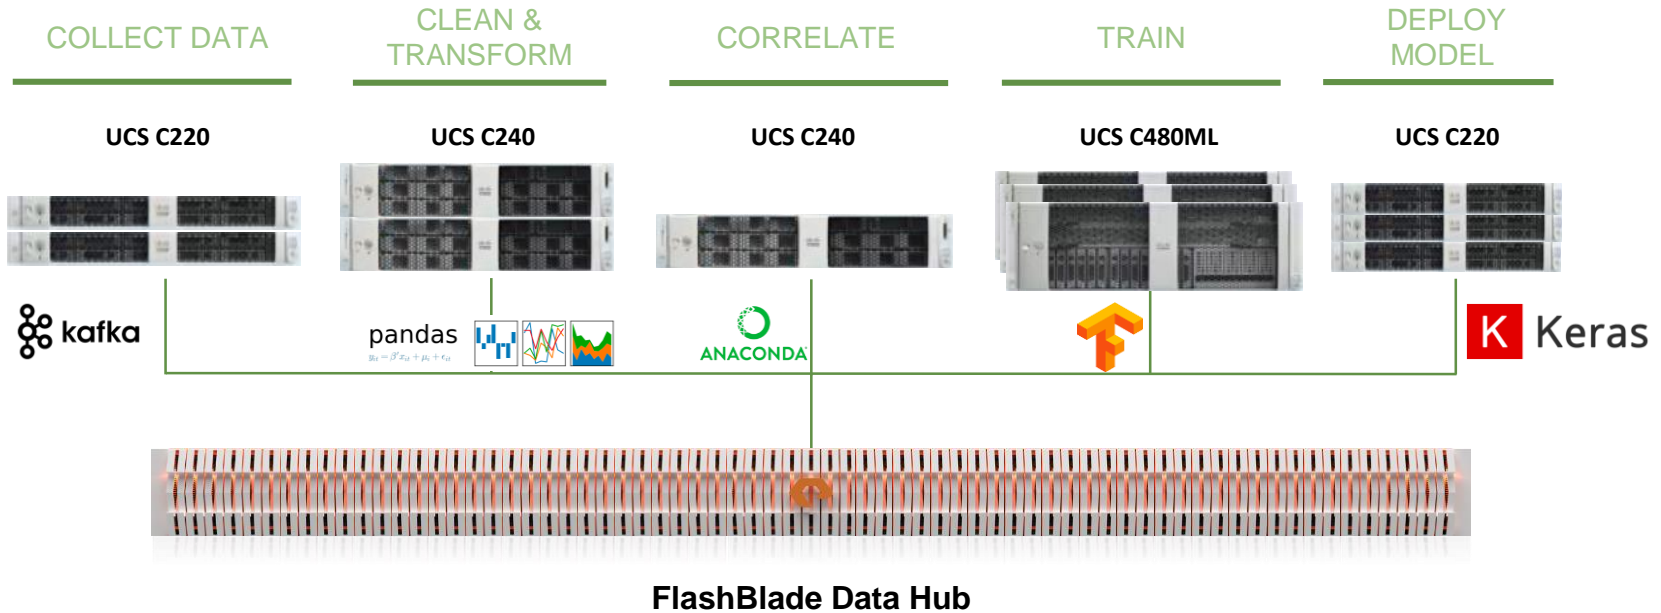


*Enabling self-service  
for  
faster iterations  
and  
experimentation.*

# SIMPLIFY YOUR DATA PIPELINE



## POWERED BY FLASHSTACK FOR AI



# SEAMLESSLY EXPAND YOUR BIG DATA ECOSYSTEM TO AI



## Cisco Validated Designs for FlashStack for AI




### Cloudera Data Science Work Bench

Hadoop Coupled with GPU Nodes for Deep Learning with Jupyter Notebook



### Kubeflow

Portable, Scalable ML Stack Enabling Rapid Development and Deployment



### Hadoop 3.1 Data Lake

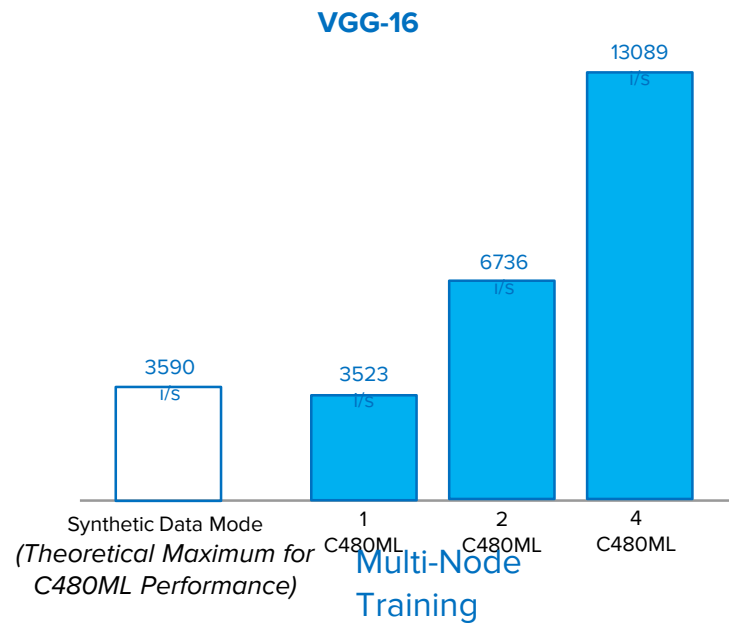
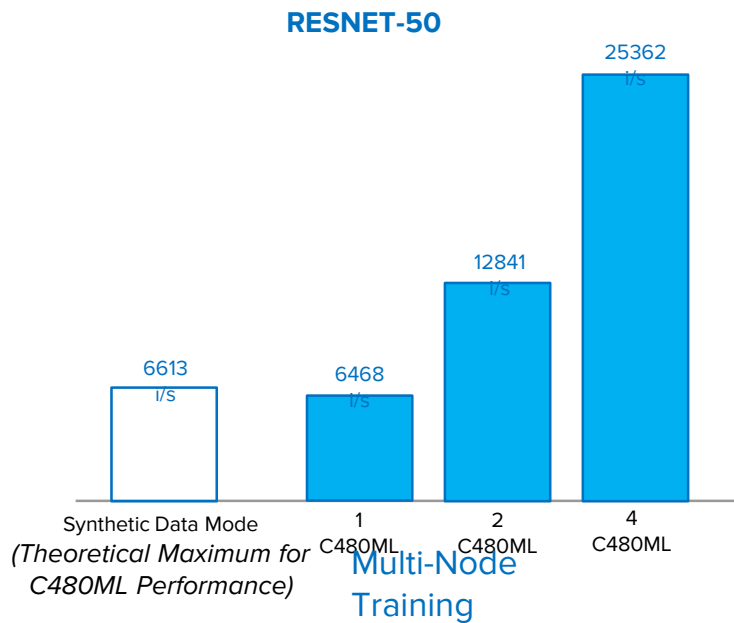
Integrate Hadoop and AI/ML: YARN Scheduling CPU and GPU with Docker Application Support



# SEAMLESS SCALABILITY FOR IT TEAMS



## MAXIMUM PRODUCTIVITY FOR DATA TEAMS



# ANNOUNCING FLASHSTACK FOR AI



## Cisco UCS C480 M5 ML: Next-Gen AI Server

- Purpose-built rack-scale solution for Deep Learning
- Eight NVIDIA Tesla V100s with NVIDIA NVLink Interconnect
- Scalable AI frameworks out-of-box

## Cisco Fabric: State-of-the-art Data Center Fabric Switching

- Massive Network bandwidth for fast data ingest with Nexus switches- 100Gb
- Multi-protocol Storage Interconnect

## Pure FlashBlade Storage: Fast, Reliable and Shared Storage

- Next-Gen all-flash storage architected for modern analytics & AI
- Seamless data sharing across the data pipeline with Pure Data Hub
- Elastically scale-out to Petabytes as data grows with zero-

*Simplify Your ML and DL Data Pipeline!*

# FLASHSTACK FOR AI



**Next-Gen Stack**  
Platform for AI/ML Workloads



**Built by Industry Innovators**  
AI/ML in a Box

**Next Gen Solutions for Next Gen Opportunities**



**FlashStack Built for AI/ML**

**For Modern Workloads**  
Ingest, Analyze, Access at Massive Scale



