servicenow

Enterprise AI

The importance of text

Debu Chatterjee

Sr. Director,

AI and Analytics Engineering

Who are we? Why are we here?

Our company - ServiceNow

We make the world of work, work better for people.

Works for you

NYSE: NOW

~6,700 Global Employees across 27 Countries

9 Pairs of Data Centers

Major Sites

Santa Clara, San Diego, Hyderabad, London, Amsterdam, Sydney, Seattle, Tel Aviv, and now **Chicago**

Strong Revenue Growth

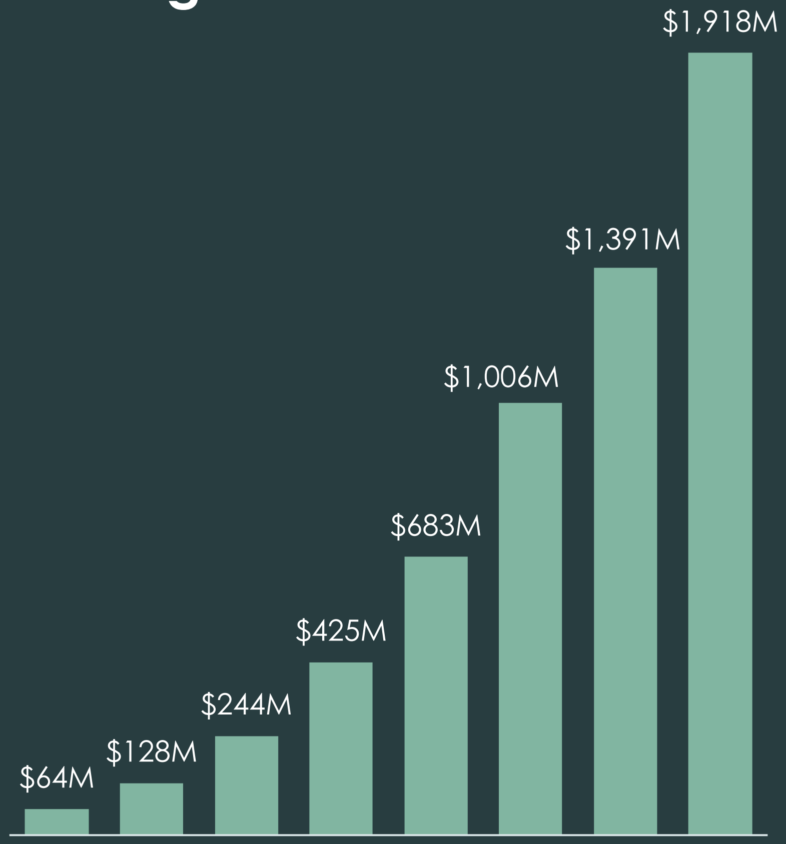

Note: 606 metrics stated for FY16 and beyond

now.

Our customers - a sampling

AIA THE REAL LIFE COMPANY
AIG
BARCLAYS
citi
CREDIT SUISSE
ING
Morgan Stanley
NYSE EURONEXT
UBS

Johnson & Johnson
Kimberly-Clark
REI
Shaw FLOORS
Steelcase
Tommy Bahama

Capgemini CONSULTING.TECHNOLOGY.OUTSOURCING
CompuCom
DELL
dimension data
EQUINIX
Hitachi Consulting
KPMG
LOGICALIS SMC Service Management Consulting
orange Business Services
SUNGARD

ASCENSION HEALTH
BlueCross BlueShield Association
BD
cdh CENTRAL DUPAGE HOSPITAL
CHRISTUS Health
FAIRVIEW
HealthPartners
NHS
THE OHIO STATE UNIVERSITY Medical Center

AMD
DOLBY
FLEXTRONICS
intel
JUNIPER NETWORKS
LEXMARK
NetApp
QUALCOMM
redhat

now

Great enduring companies

They are
purpose driven

Innovation <u>and</u>
execution

They invest in
talent

Exhibit "will to
fight"
and "will to win"

now.™

Innovation - the company we keep, and we are #1

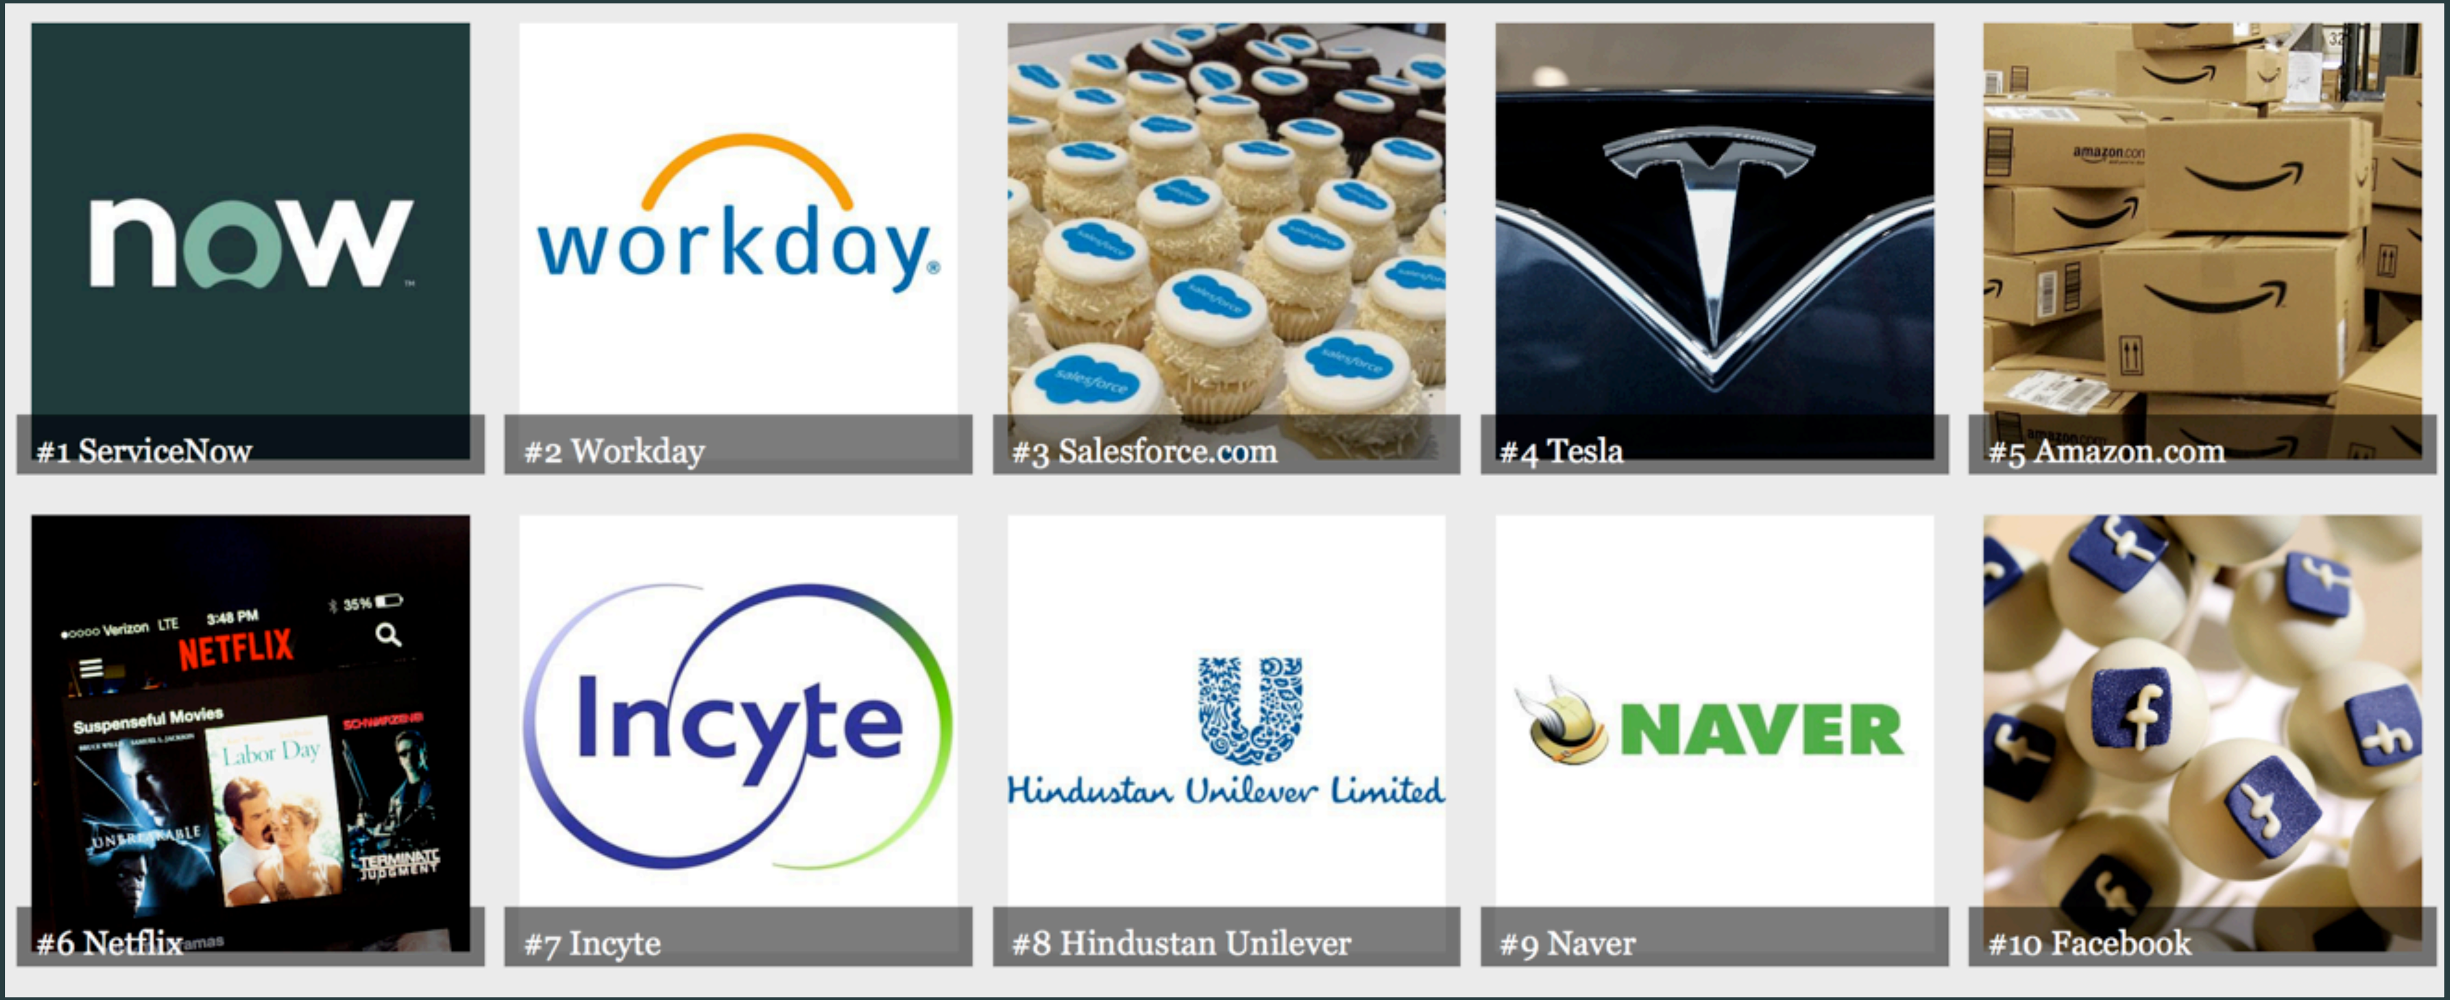

Priority: Develop talent - we are hiring in Chicago

Machine learning engineers

UI/UX engineers

All levels, and experience

What do we “Product” folks do?

Why AI using unstructured textual information is important

Emerging products

- Security
- Customer Service
- HR

IT products

- ITSM
- ITOM
- ITBM & ITAM

Now Platform™

- User and Service Experience
- Service Intelligence

The Now Platform scale

Production instances

~7,500

Customer transactions per month

~12 Billion

licensed users

~18 Million

# of Cloud API calls per month

~18 Billion

Native mobile applications

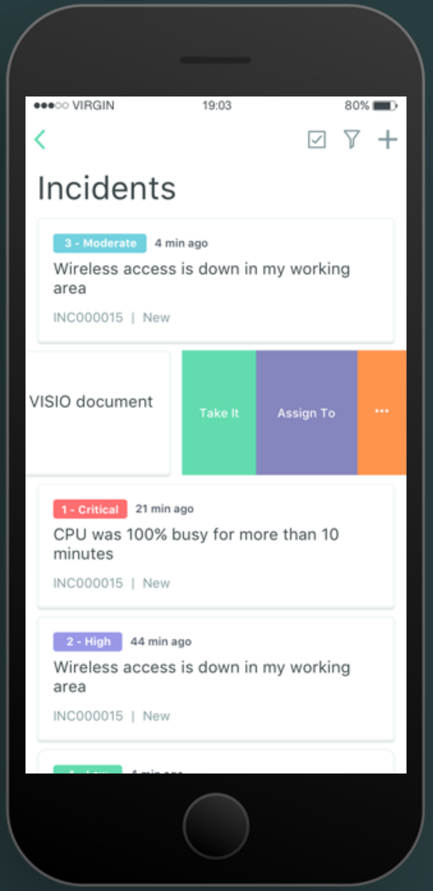

Incidents

Swipe Actions

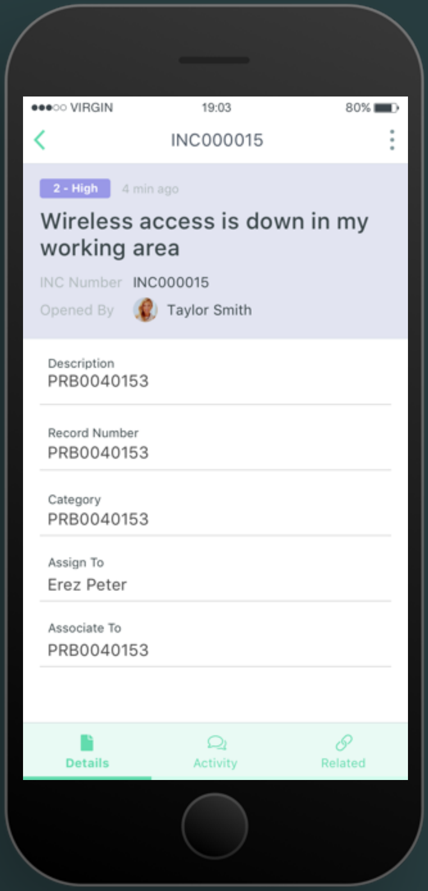

Incidents Details

Call/Text/GPS

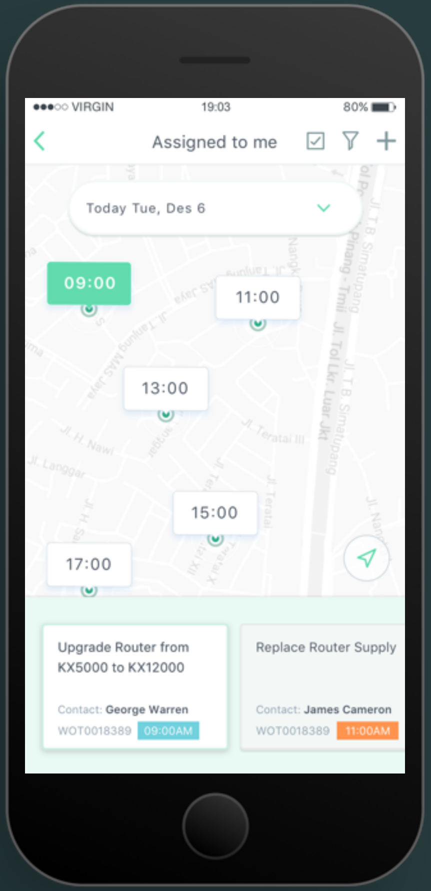

Map View

Filter by Time/Location

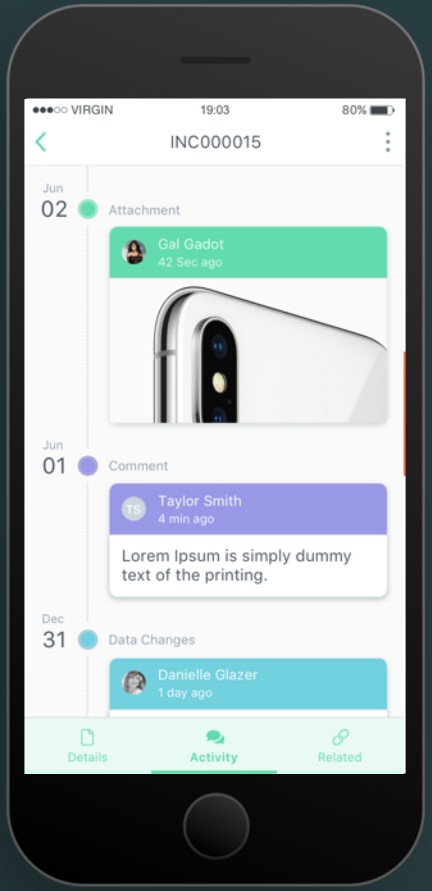

Activity Stream

Add Attachments, Images, Comments

Virtual Agent

now™

Agent Workspace

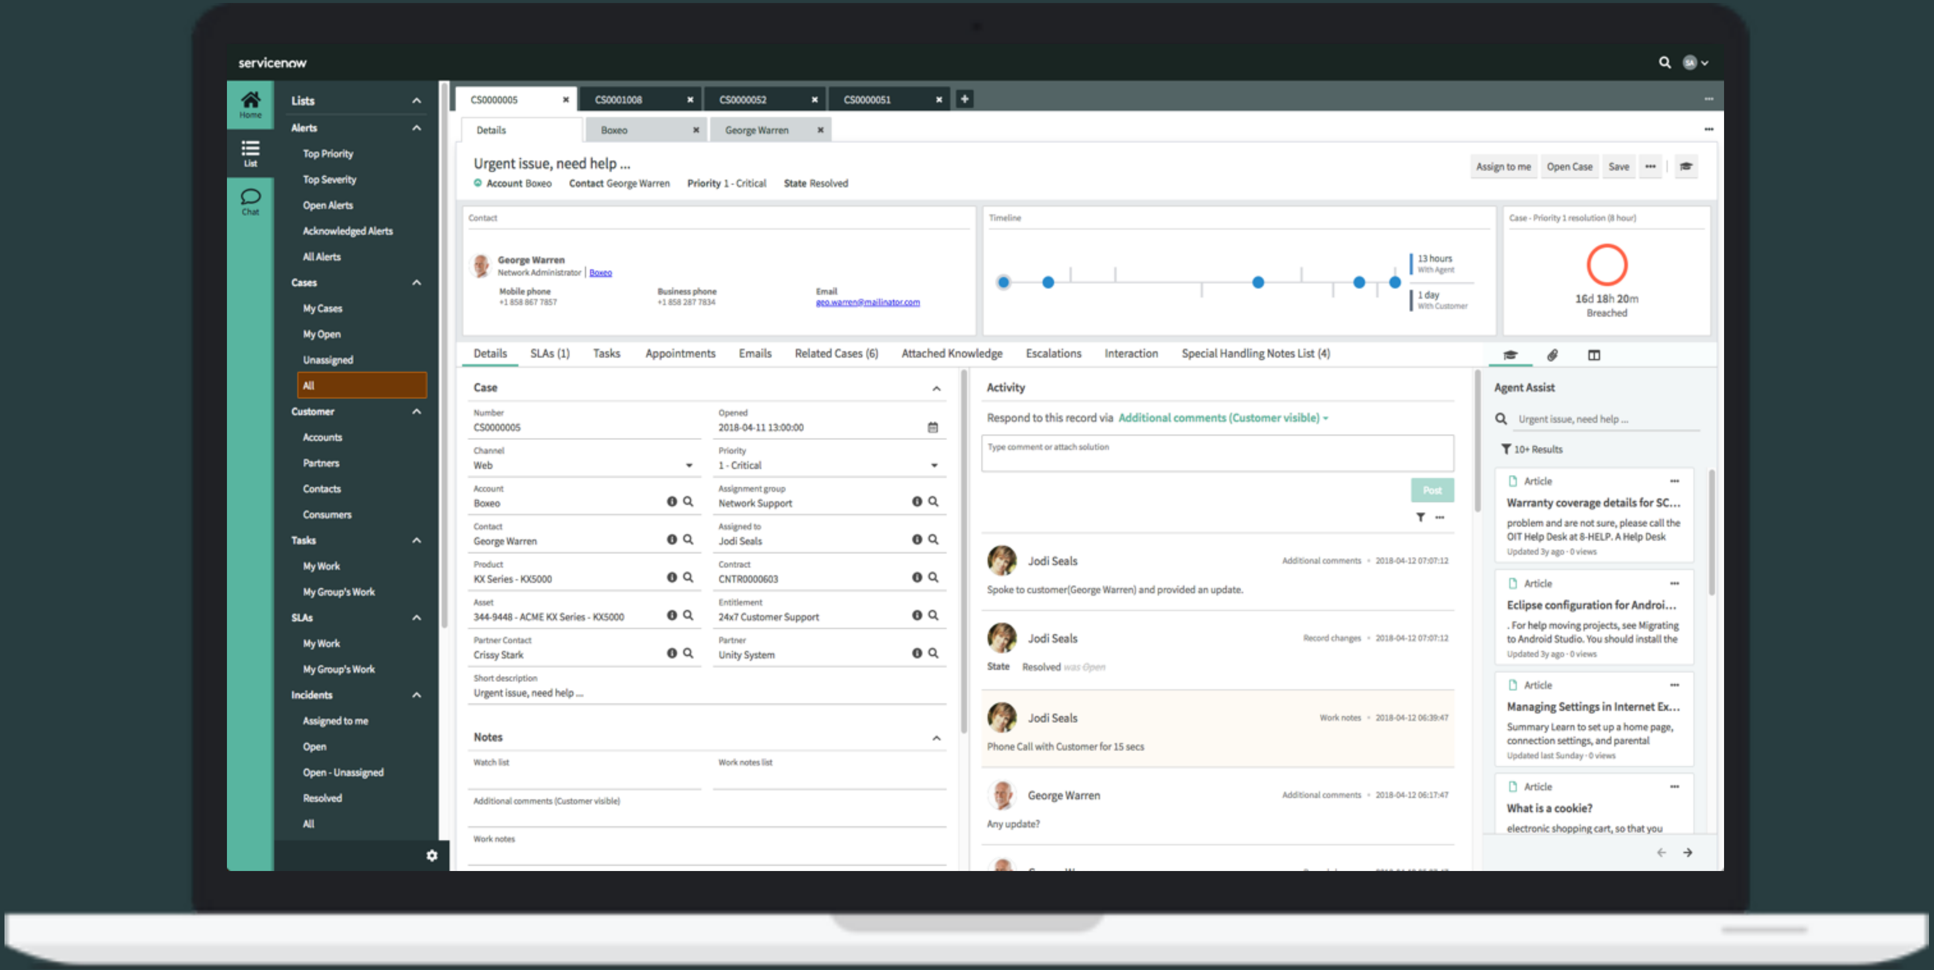

now

Opportunity for AI driven service experience

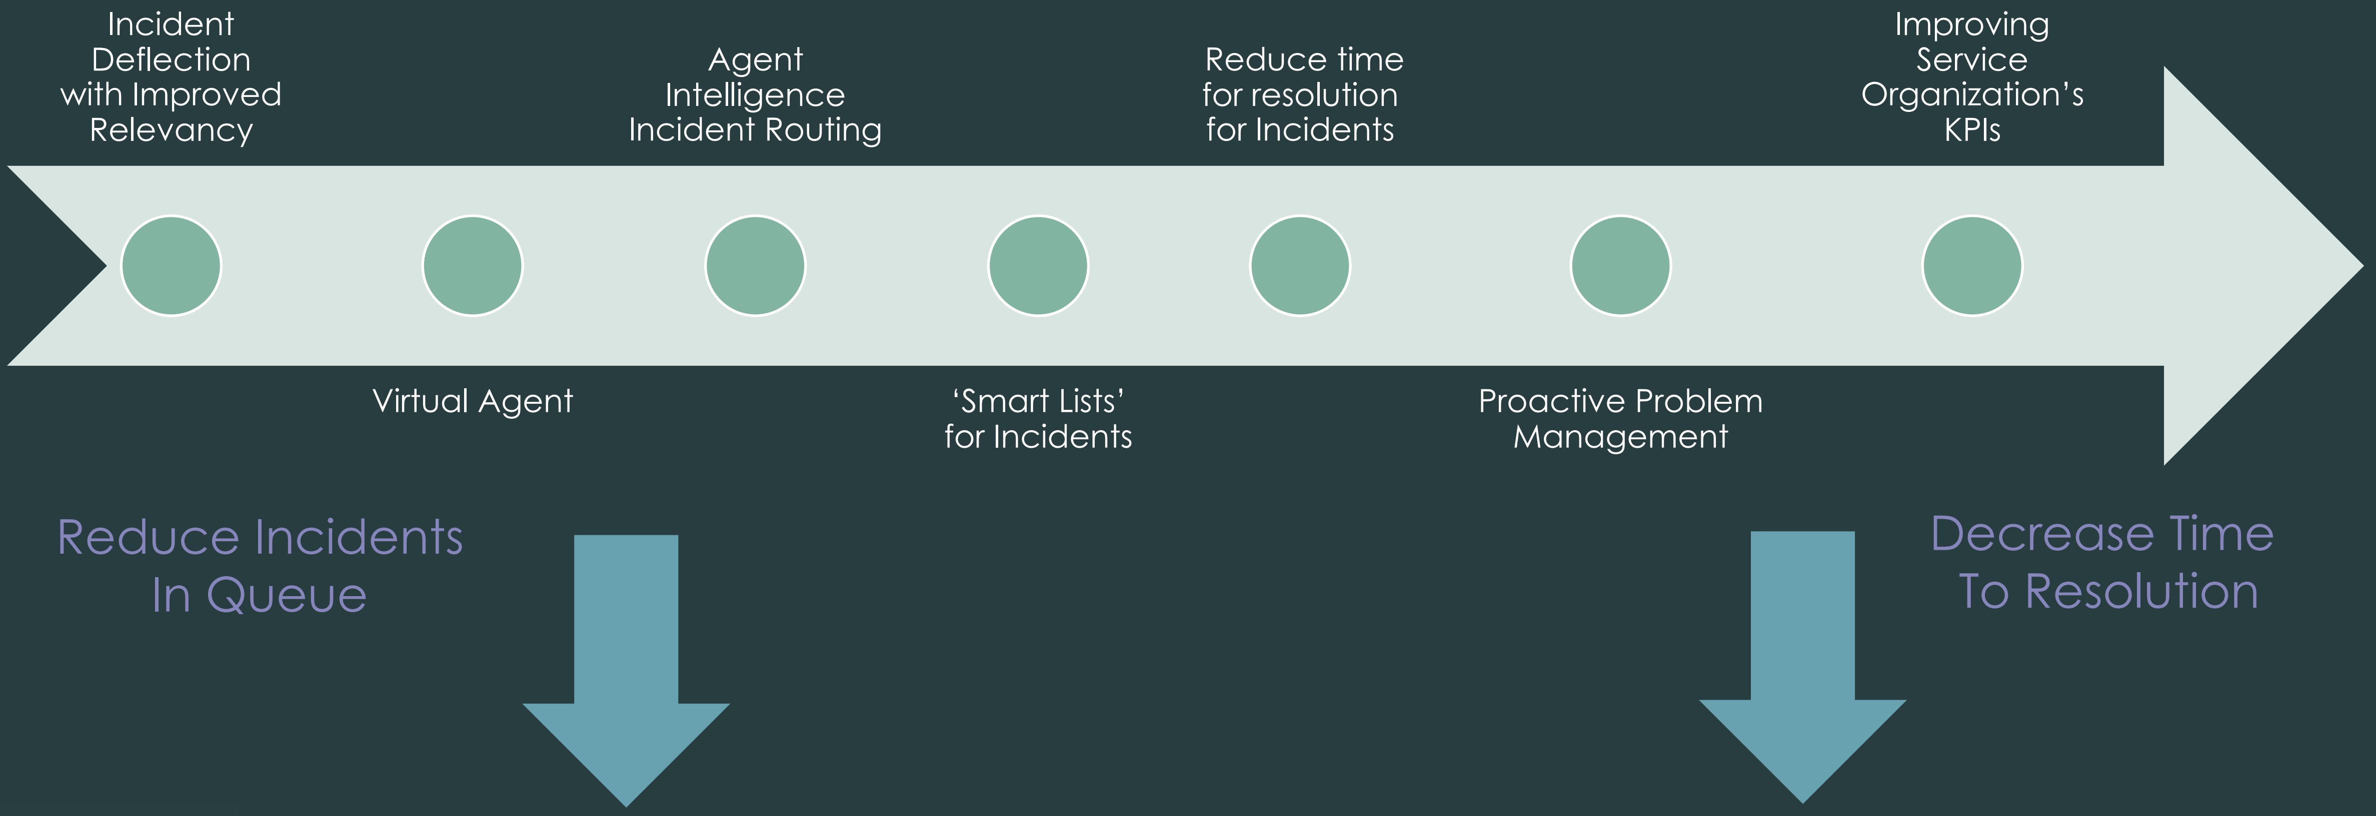

now.

An AI framework drives it all!

Ranked Search,
NLP
Classification
Association Rules

NLP
NLU
Intents
Entities
NLG
Recommendations
Sentiments

Classification
Regression
Forecasting

Clustering
Similarity
Scoring

Forecasting
Optimization

Anomaly Detection
Outliers
Recommendations
Association Rules
Classification

Prescriptive
Explainable AI

Hence, the importance of AI using text in the enterprise

- **Enterprise** very different from **Consumer** (ServiceNow data ≠ **Twitter** data)
  - Consumer apps have **one** big **database** of text to study and build models from
  - Enterprise has security rules: we have **5k+** customers all in **separate instances**
  - Challenge is to **build many** similar **models** on **different datasets** with expectation that they work **OOB**
- Language "at work" is different from language "at home"
  - Word embedding (eg from Google) on "home" language don't capture work jargon (NER is hard)
  - **URLs**, file directory **paths**, product **versions** – all potentially meaningful, but **inflate vocabularies**
  - "Sentiment analysis" research not on point. "**propensity to escalate**" matters rather than "**negative**"
  - **Many** contemporary chatbots are working to a "keep the conversation going" metric, **longer is better**
  - **Customers** seeking support/answers **want FACTS** and **ACTION**, not chit-chat. **Shorter** is better.
  - I can ask IT for a new laptop, but am I allowed to have one? Can we **blend rules/logic with NLP** effectively?
- We cover multiple **domains** (HR vs IT), and many '**subdomains**' (IT but at different customers)
  - **Transfer** learning or **meta-learning** techniques should be relevant
  - Learn **inductive biases** to pick up customer's jargon. **We know it's an IT shop**, just learn local vocabulary.
  - Combine Knowledge Base + NLP? **NER** plus **relations**: team X in org Y working on project W for release Z

Volume and Variety

Incidents: short to medium text

Customer	Yearly Volume
A	673K
B	426K
C	187K
D	187K
E	171K

Alerts: short to medium text

Customer	Yearly Volume
A	9.6MM
B	1.8MM
C	380K

Knowledge Base: medium to large text

Customer	Yearly Volume
A	160K
B	131K

now

Class of scaling problems

- Effective methodologies for combining hierarchical data
  - Incident records have data at different levels
    - Description is a low-level feature.
    - Category or other tags is a higher level feature.
  - **Need to combine high and low level features.**
- Text matching under varying text lengths
  - A new incident may have a pre-existing resolution in an article in the knowledge base. How to find the best one?
    - Incident description is a short text
    - KB articles are long texts.
  - **New methodologies needed for short-to-long text matching.**
- Vector representations of text for deep learning models
  - Text data can be represented many different ways: TF/IDF vectors, paragraph vectors, word vector based convolution, etc.
  - **Determining the best representation.**
  - **Combine generic word2vec/glove vectors and customer/domain specific vectors to get higher precision**
  - **Get these things to work across languages scalably**

Challenges with ML/AI solutions

- Encapsulation is more difficult in ML solutions
  - Changing anything changes everything (CACE)
- Data dependencies more severe than logic dependency
- Common anti-patterns
  - Real systems = 5% ML Code + 95% “glue” code
  - Pipeline jungles
- Testing is much more complex
  - Exact reproducibility, the backbone of traditional testing can be often compromised
- Multiple Languages as in English, French etc
- Enterprise Specific Language

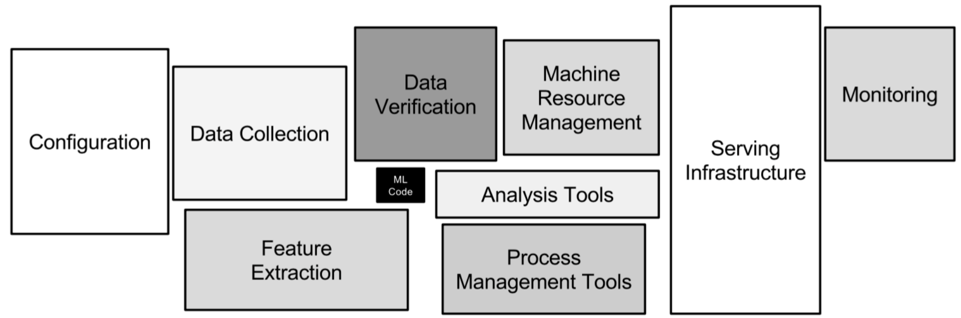

Figure 1: Only a small fraction of real-world ML systems is composed of the ML code, as shown by the small black box in the middle. The required surrounding infrastructure is vast and complex.

Source: Sculley, et al NIPS 2015

Many other challenges remain

Need too many ’labeled’ examples.

How does the math (SGD) work here?

Many other challenges remain

Easy to manipulate models

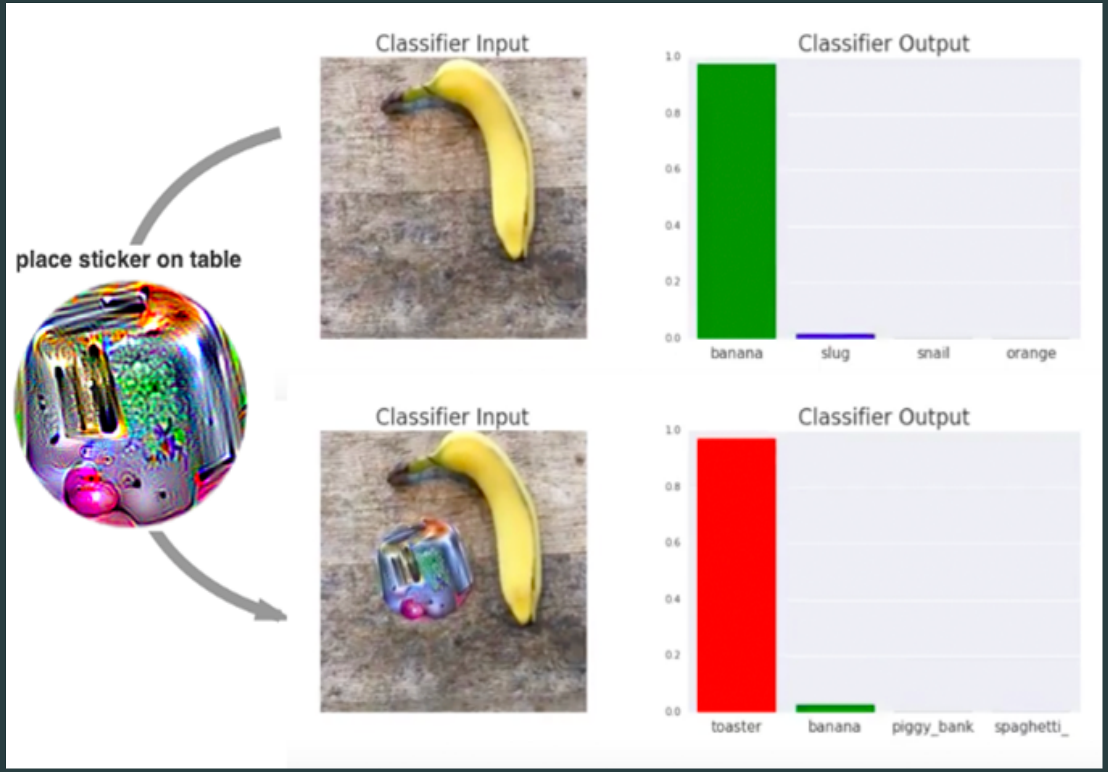

Difficult to explain

now.

Why ServiceNow!

AI is the key differentiator

Enterprise AI is lot about text

And, we have fun being a guiding hand to each customer one at a time

now™

Thank you!