Cloud-based IT Log Analytics

Christian Beedgen

Kumar Saurabh

Agenda

Overview

Team

Market Size

Problem Statement

The Next Generation

Differentiators

Competition

Go To Market

Economics

Roadmap

Summary

Overview

Cloud-based IT Log Analytics

Service to manage and analyze IT logs

$2.5 Billion market size

Current products have high TCO, are services-heavy

Easy to get started, lower TCO, superior intelligence

Team of log management veterans, to be completed

Series A – customer-focused development process

Team

Christian Beedgen

ArcSight since 2001, Chief Architect, Director of Engineering

Lead ESM server developer

Built ESM server team, managing 20 people in server and UI teams

Named on 2 granted patents, 7 patent applications in process

Past experience at **Amazon**, **Gigaton**, **Cleverlearn**

Kumar Saurabh

Data Architect at **Mint.com**

Single handedly built Mint's data analysis infrastructure

ArcSight 2001-2008, Director of Engineering, managing 12 people

Lead for Analytics and Solutions Team

Named on 2 granted patents, 2 patent applications in process

Market Size ~$2.5 Billion

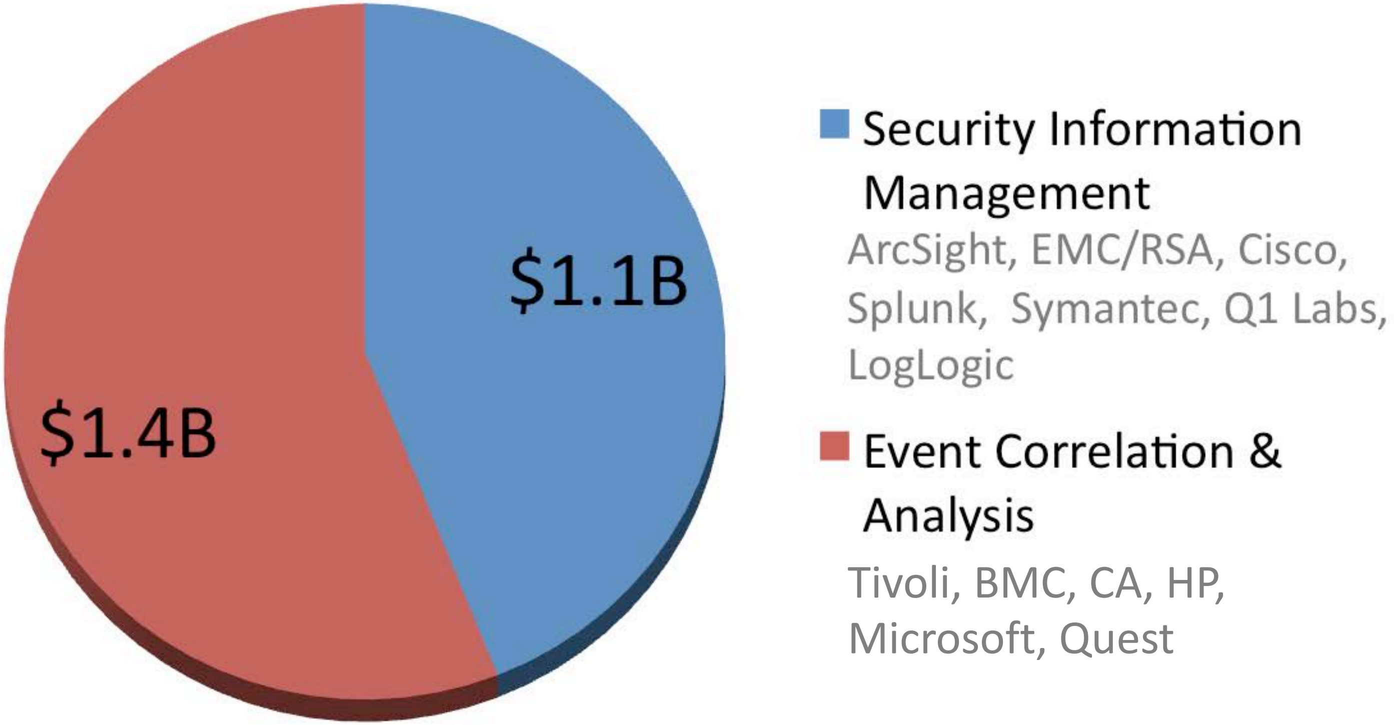

Source : Gartner/Dataquest

Key Drivers: Compliance, Security, Operations

Key Drivers

Compliance is not optional

"What is the primary motivation for adopting or using security information management (SIM) within your enterprise?"

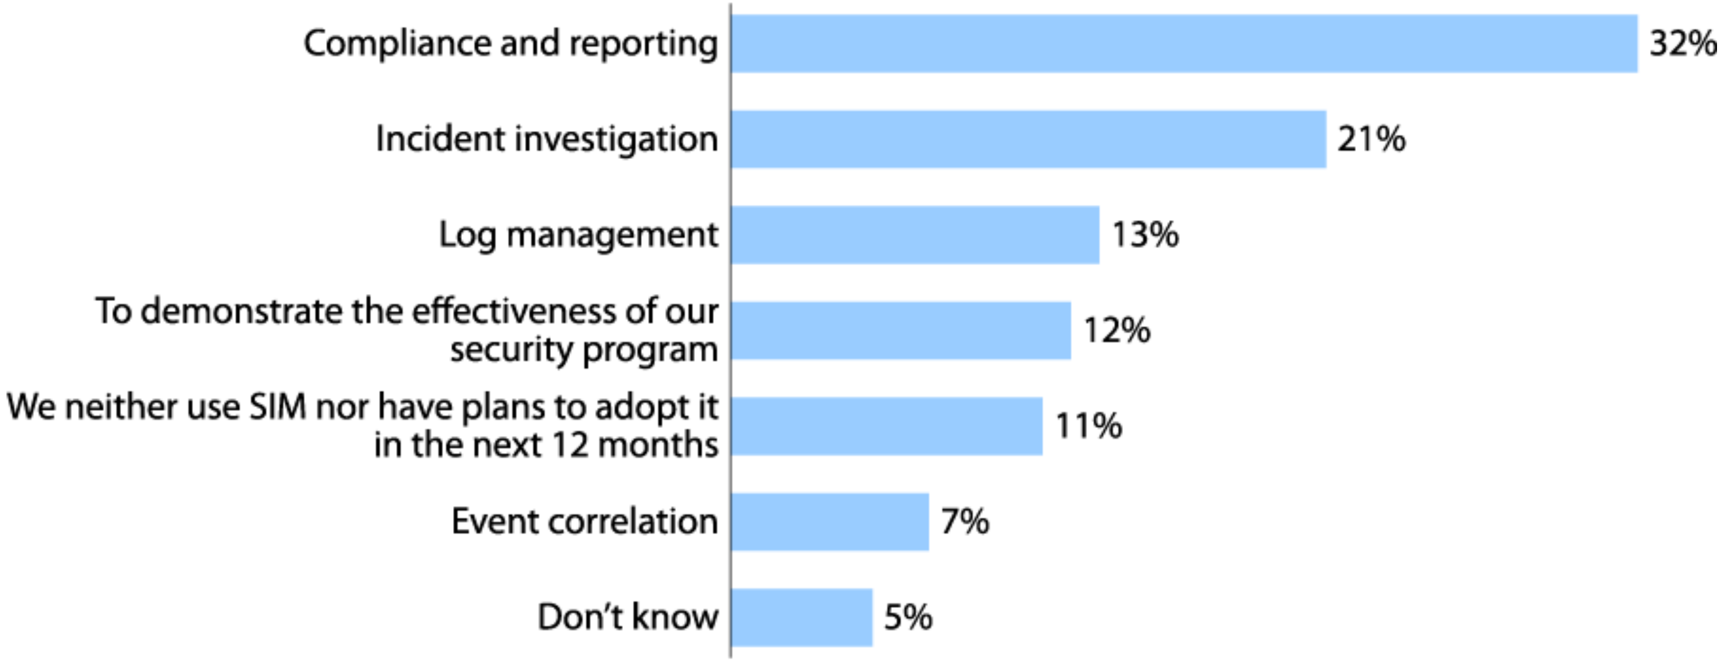

Base: 1,335 North American and European enterprise and SMB security decision-makers who expressed interest in adopting SIM
(percentages do not total 100 because of rounding)

Source: Enterprise And SMB Security Survey, North America And Europe, Q3 2008

FORRESTER® April 2009 **"Market Overview: Security Information Management (SIM)"**

Problem Statement

Today’s market leading products are:

Premise-based

Enterprise sales cycles, installation and upgrade hassles, expensive hardware, DBAs, sysadmins required

Not scalable

Not inherently clustered, scaling introduces tradeoffs and data fragmentation

Challenged with log parsing

Either simply don’t parse or require parsing at collection time, need constant software upgrades

Not context-aware

Identities, network assets, service dependencies are all critical for correlation and prioritization

Customers operate in silos

Insight gathered by one customer is hard to share; no cross-customer data mining

Not community-aware

Exchanging of solutions is a manual process, there’s no marketplace

The Next Generation

1 **Cloud-based service**

Easy sale, quick delivery, ongoing upgrades, no care and feeding

2 **Seamless scalability**

Built from scratch for big data, leverages large-scale processing

3 **Machine-driven log parsing**

Extracting structure from raw logs is foundation for analytics

4 **Context modeling**

Logs need to be analyzed in their real world environment

5 **Global IT log intelligence**

Data mining leads to insight shareable across all customers

6 **Built-in community**

Not everybody is an expert, and even experts exchange findings

Deliver superior log management for compliance, security and operations in a scalable, easy-to-adopt cloud-based service

Target Market

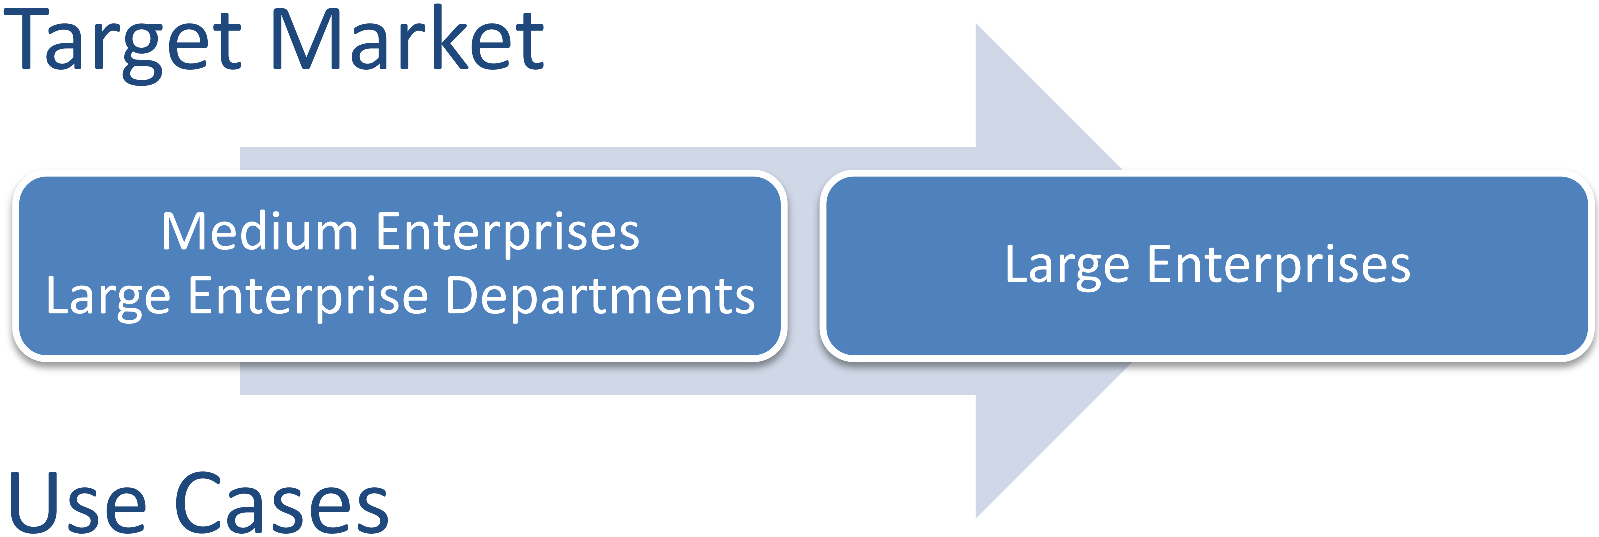

Medium Enterprises
Large Enterprise Departments

Large Enterprises

Use Cases

Compliance

- PCI, SOX, HIPAA, NERC
- Log Retention & Review
- User & Resource Access

Security

- Incident Response
- Data Protection
- Threat Intelligence

Operations

- Troubleshooting
- Business Continuity
- Service Levels

High-level Solutions Architecture

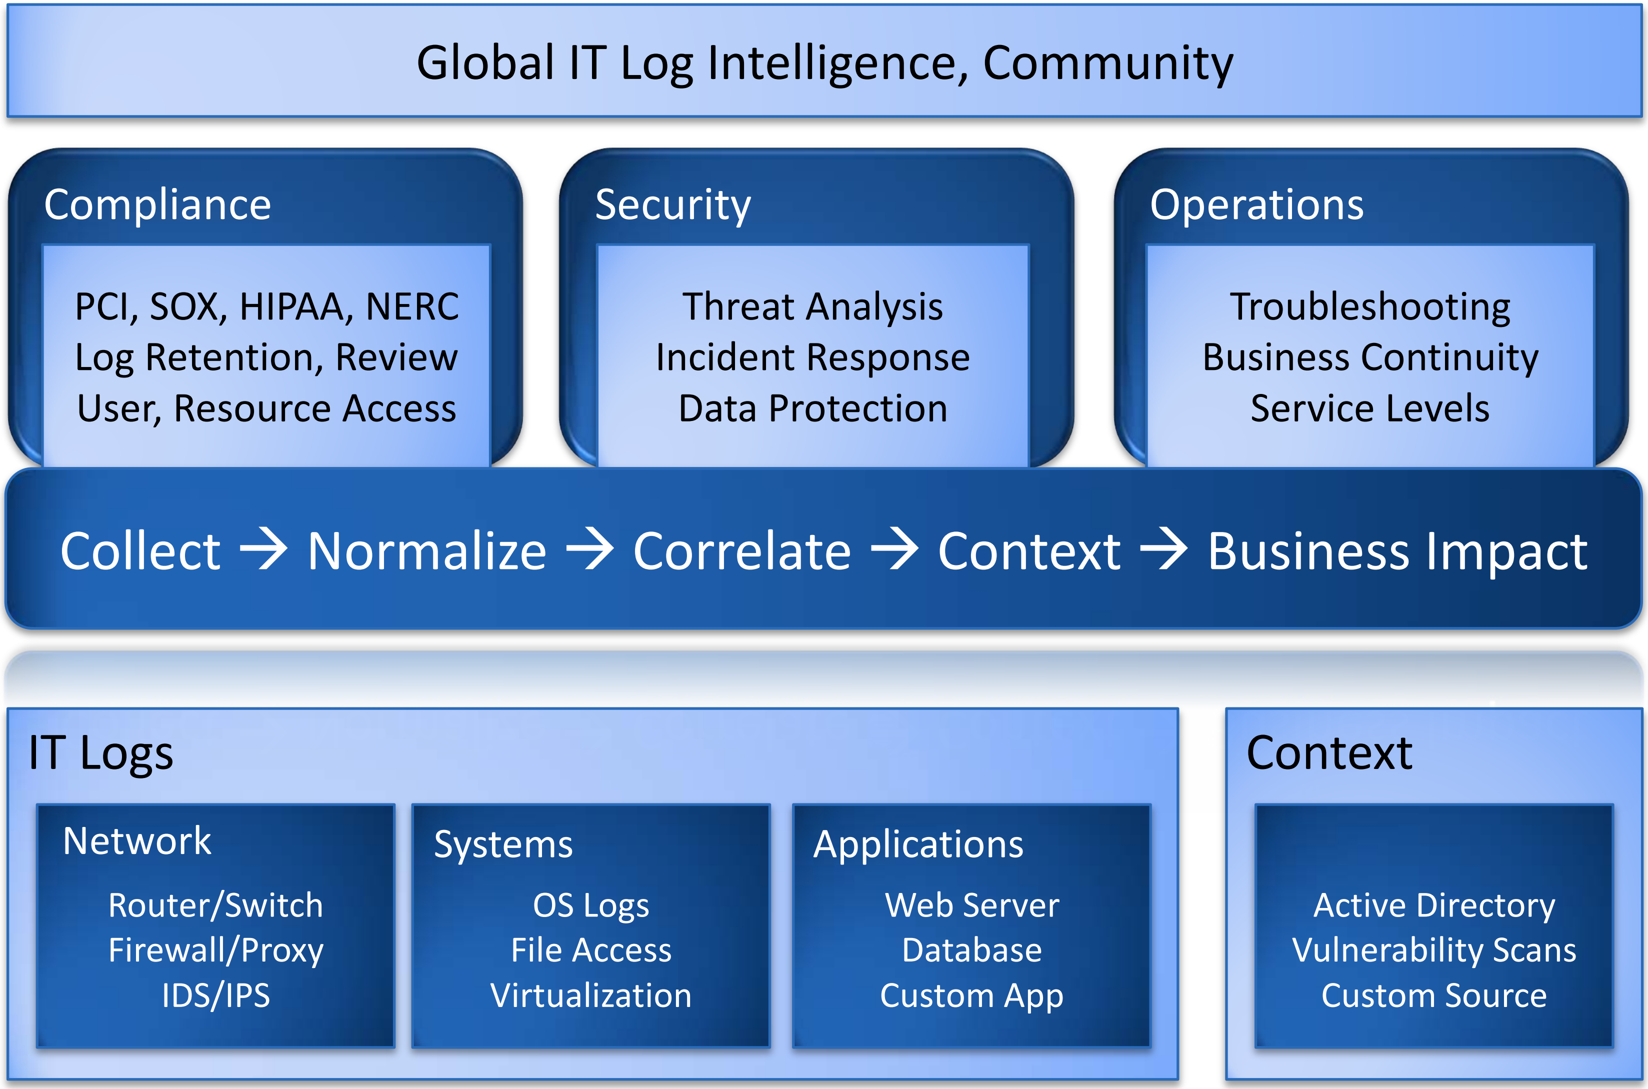

Global IT Log Intelligence, Community

Compliance

PCI, SOX, HIPAA, NERC
Log Retention, Review
User, Resource Access

Security

Threat Analysis
Incident Response
Data Protection

Operations

Troubleshooting
Business Continuity
Service Levels

Collect → Normalize → Correlate → Context → Business Impact

IT Logs

Network

Router/Switch
Firewall/Proxy
IDS/IPS

Systems

OS Logs
File Access
Virtualization

Applications

Web Server
Database
Custom App

Context

Active Directory
Vulnerability Scans
Custom Source

High-level Platform Architecture

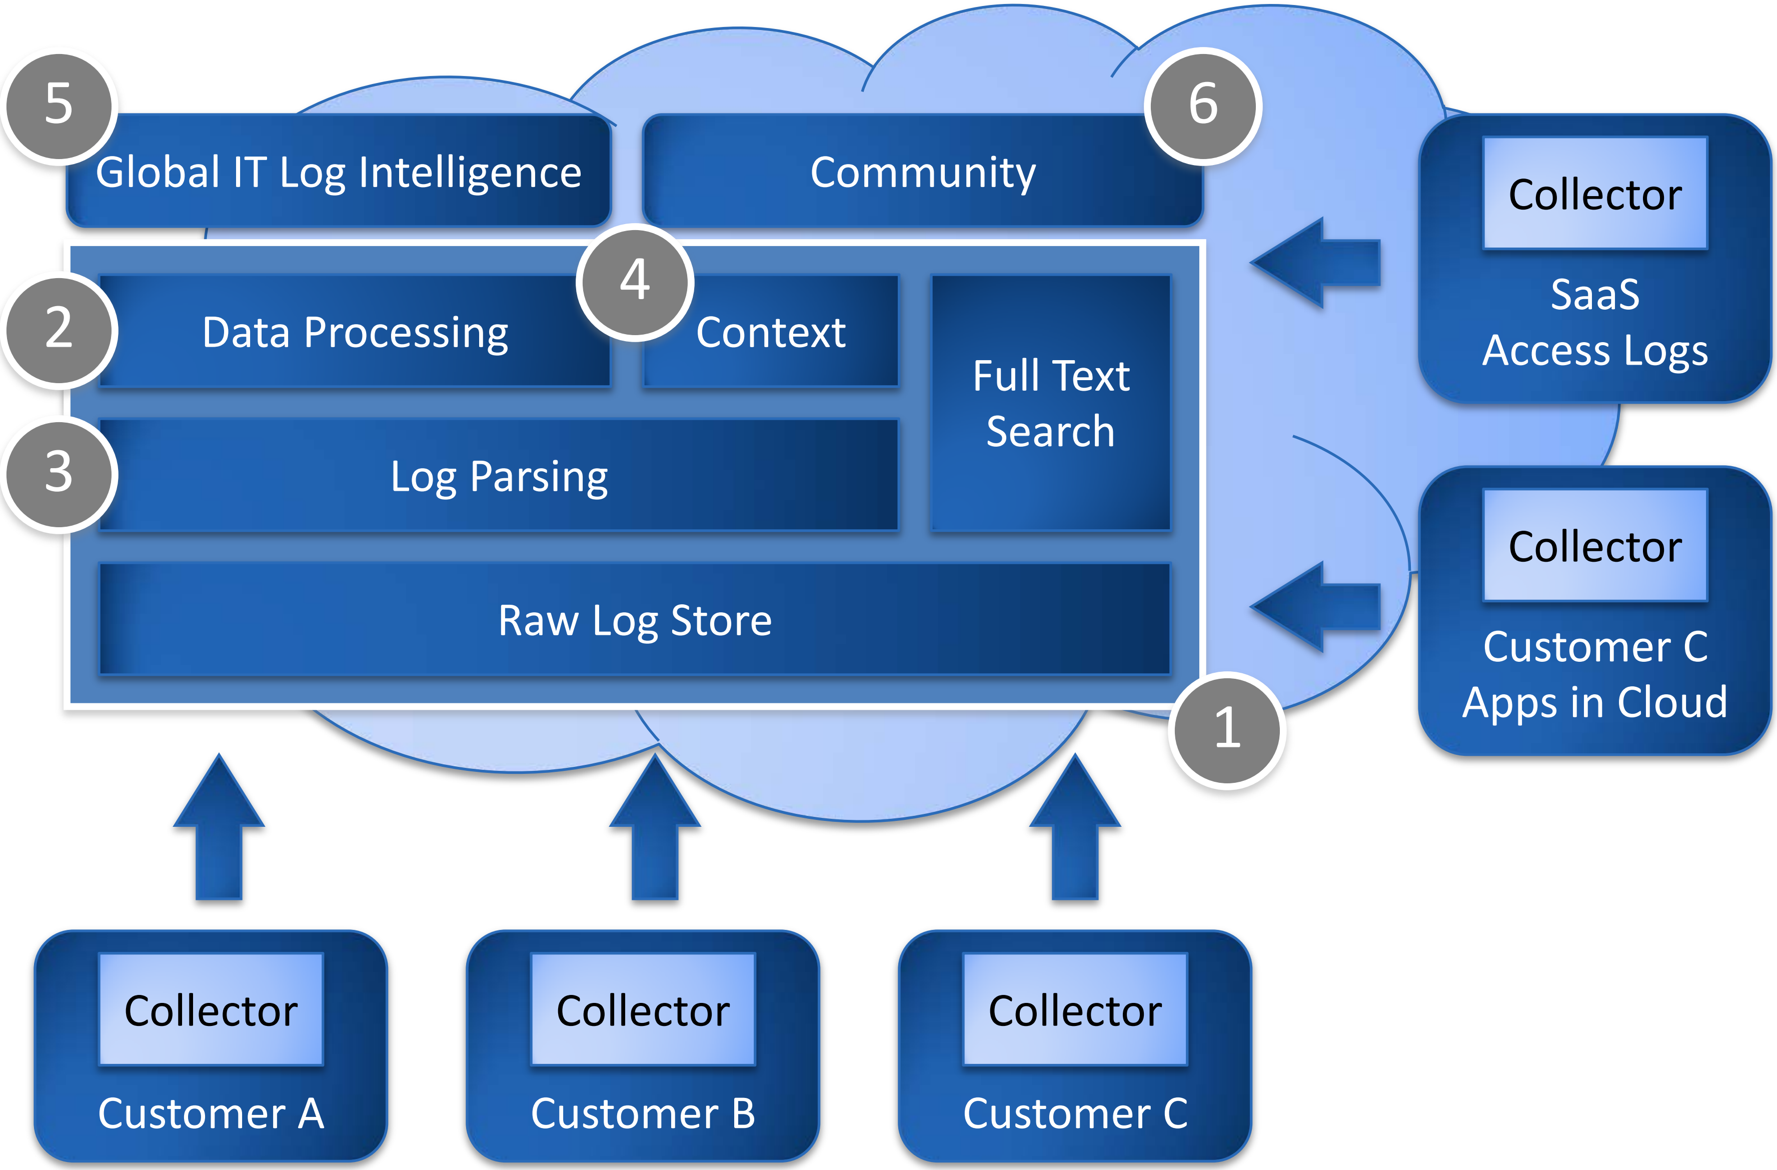

Log Management Architecture Today

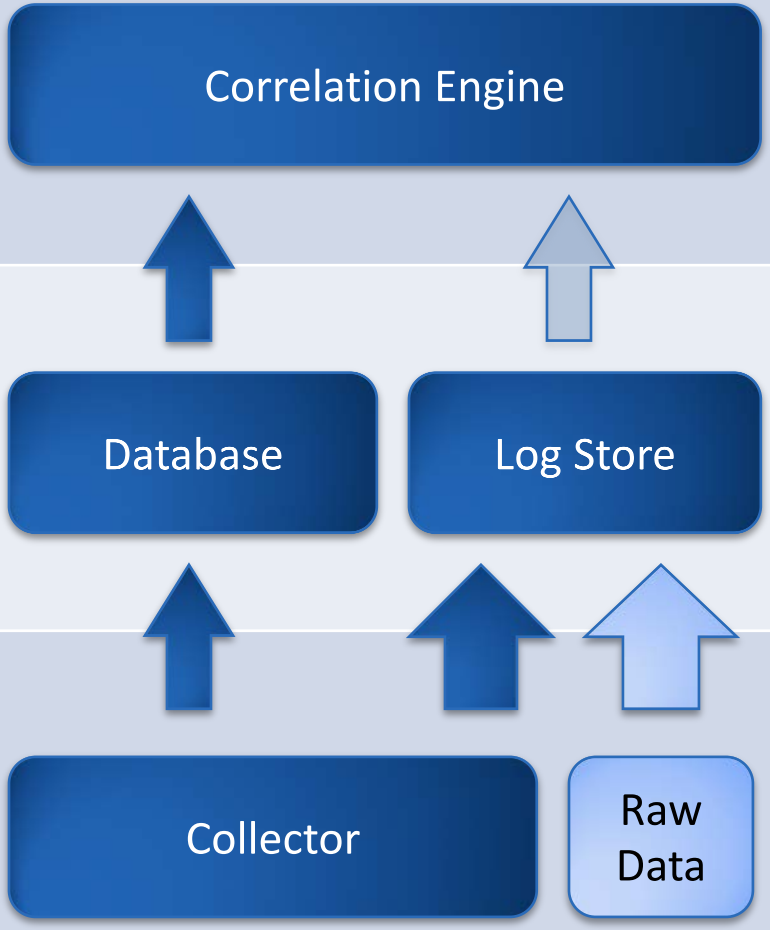

Analytics
Workflow

Compressed log store
Full-text search

Collection
Log parsing

Log Management Architecture Today

Log parsing challenges

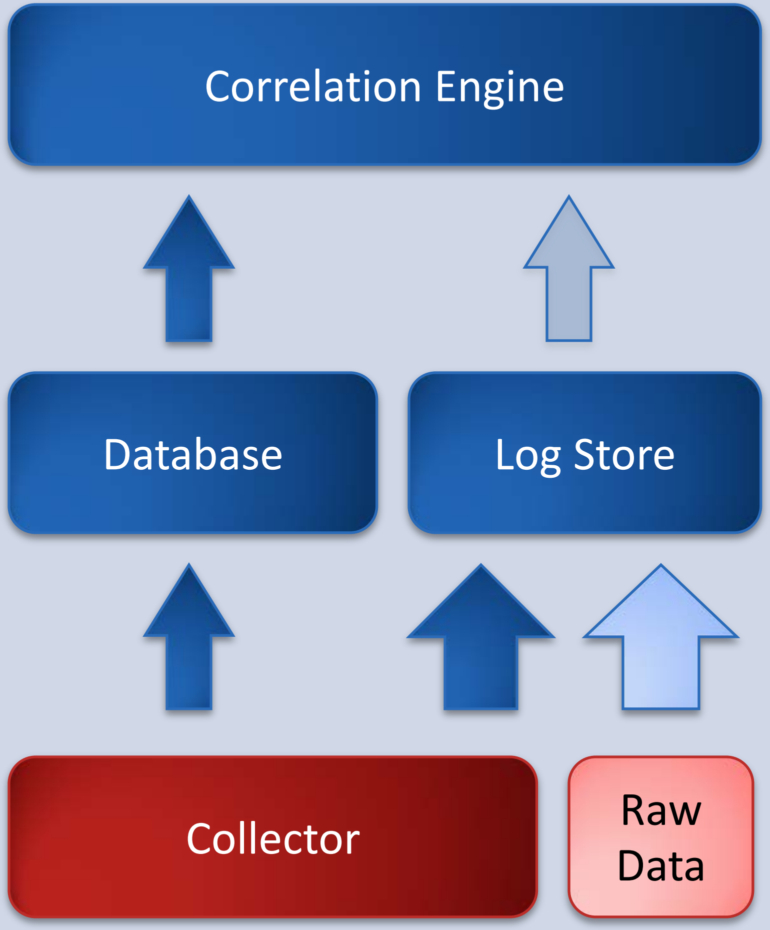

Parser creation and maintenance is manual and expensive

Limited vendor support for custom log sources like applications

Data is often not parsed, or parsed incorrectly; custom applications ignored

Next Generation Architecture

Machine-driven Log Parsing

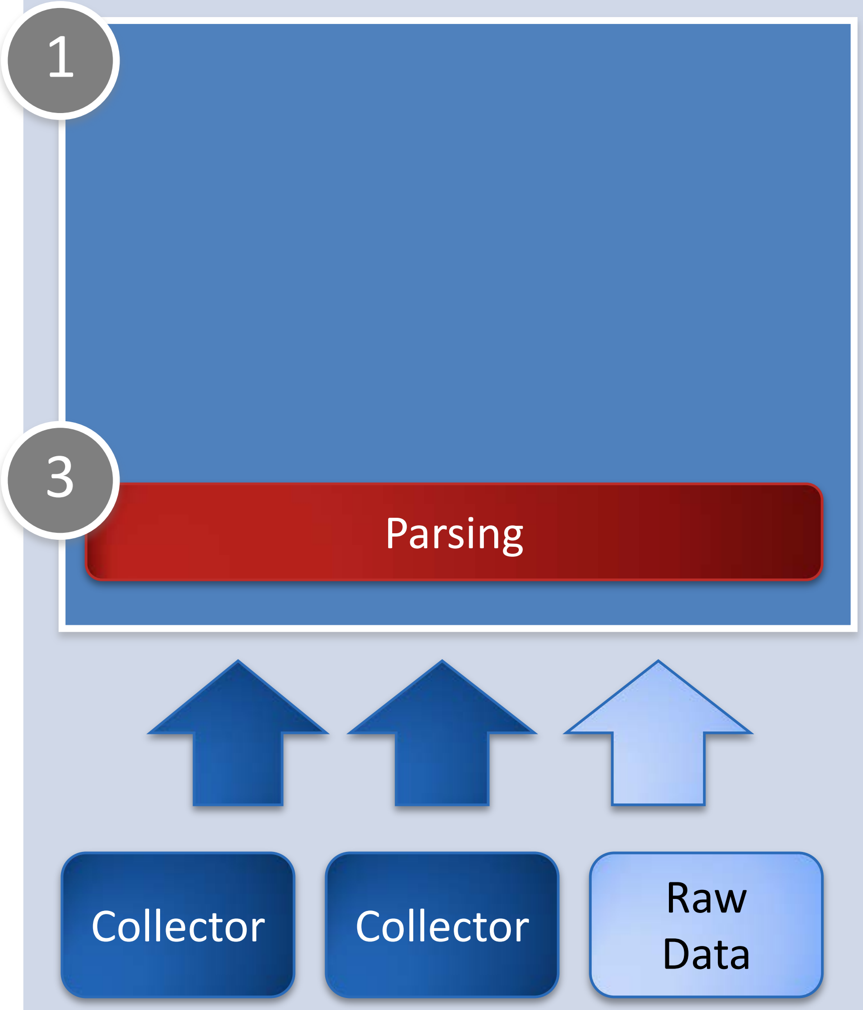

Efficient parser development driven by structure inference

Customer can add knowledge about custom sources like applications

Better parser support enables more log sources, which enables superior analytics

Log Management Architecture Today

Scalability tradeoffs and data fragmentation

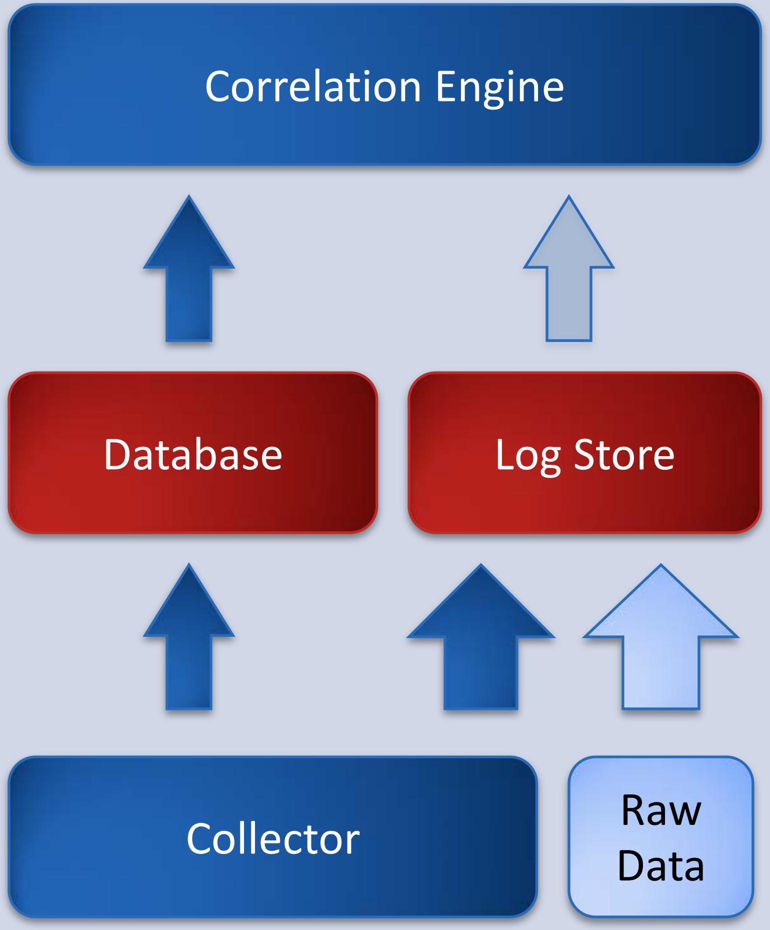

Database hardware is expensive, performance is still a major problem

Most data in the Log Store is not available for analytics

Performance or intelligence – choose one

Next Generation Architecture

Seamless scalability

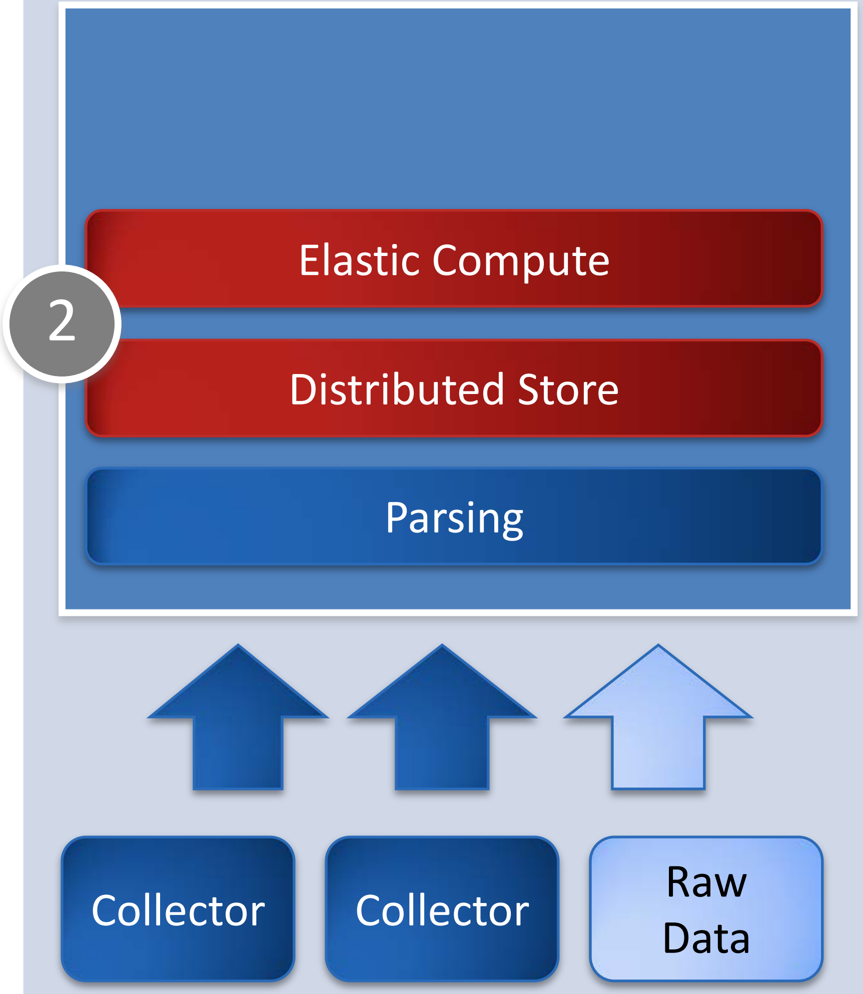

Distributed store is built for scale, no excuses

Elastic compute layer can tackle any analytics efficiently

Better intelligence and faster response times

Log Management Architecture Today

Customers operate in silos

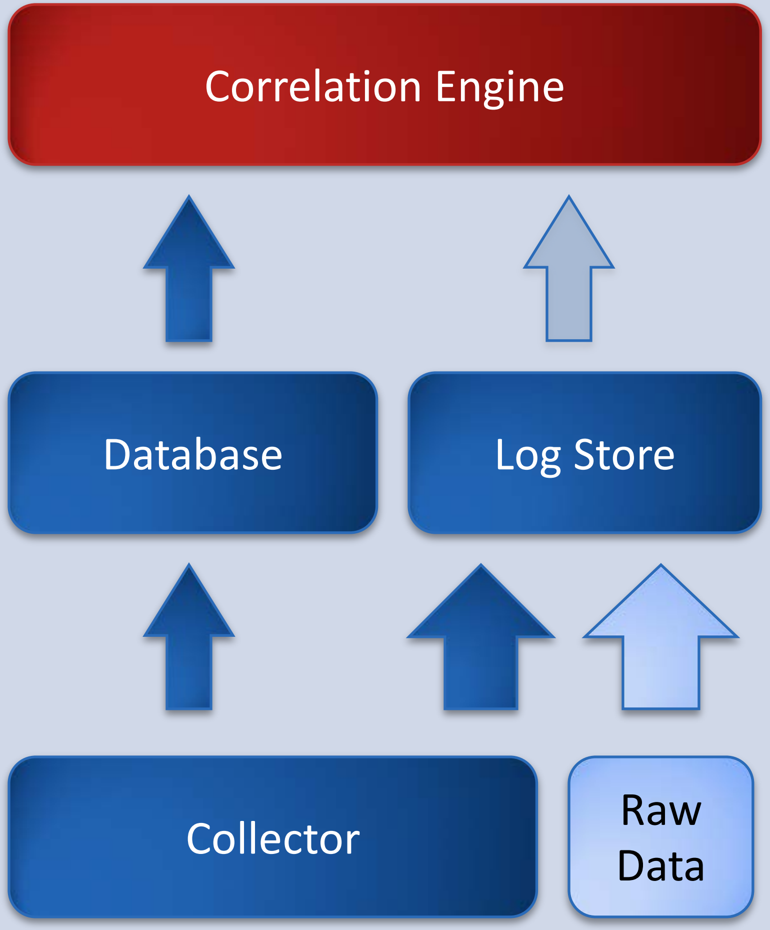

Advanced analytics cannot run due to database limits

Customers don’t always have analysts on staff, vendors don’t have the data

Incidents that could have been detected with better analytics are missed

Next Generation Architecture

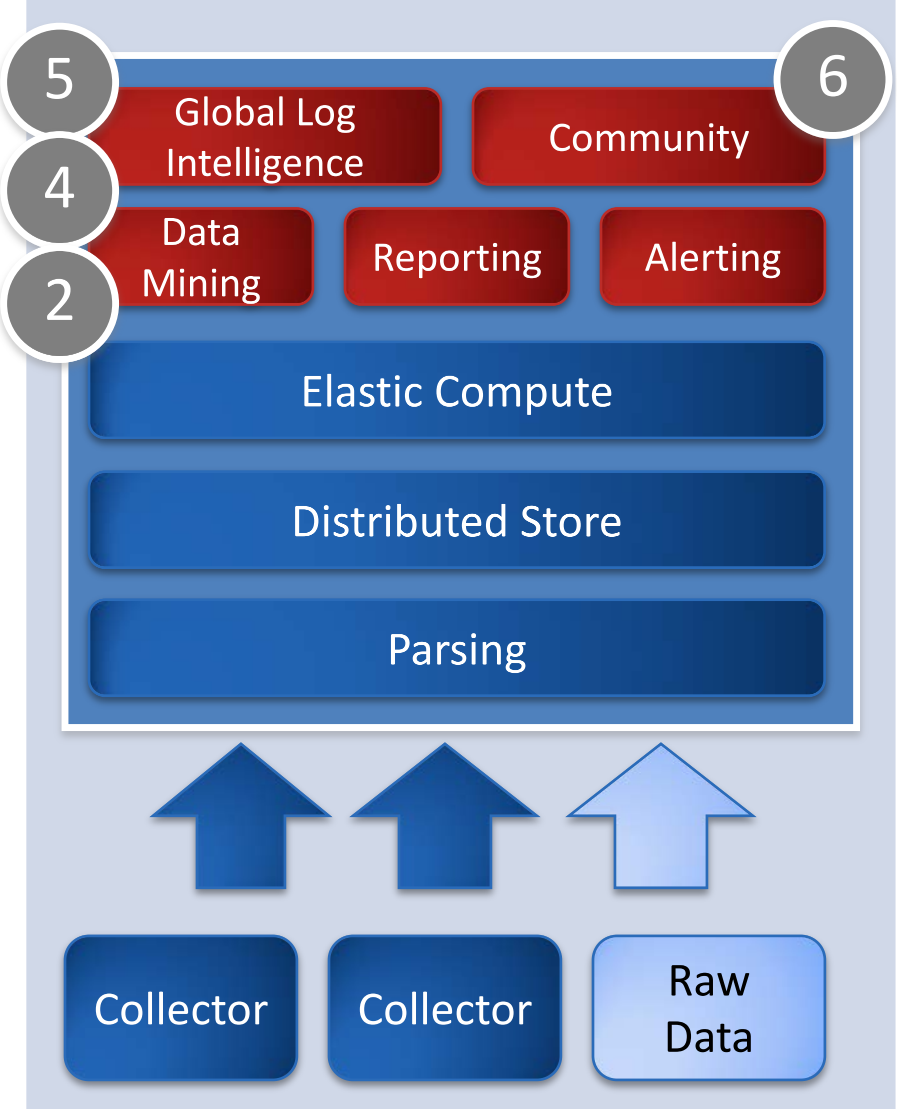

Makes it feasible to analyze all data from every customer

Enables deep analytics – root cause detection, network graph analysis, anomaly detection

Extracts and delivers actionable insights

	Differentiator	Customer Benefit
1	Cloud-based service No enterprise software release cycles	Quick initial service delivery Seamless ongoing upgrades No deployment or sizing hassles
2	Seamless, transparent scalability Headroom to deal with all the data	Instant slice and dice analytics No data fragmentation 100 EPS today, 10,000 EPS tomorrow
3	Intelligently evolving log parsing Automated structure inference	Custom app logs useful right away Structured data beats unstructured data Once seen, available everywhere
4	Context modeling Identities, network elements, services	Real world environment context for logs Business impact correlation Enables risk modeling
5	Global IT log intelligence Pattern discovery	Solution recommendation system Zero day discovery of emerging threats Defense for all customers
6	Built-in community Frictionless sharing	Expert exchange Sharing of analytics content Service → Platform

Competition

ArcSight	loglogic	splunk>	ALERTLOGIC	Cloud-based Log Analytics
Enterprise software & services, appliances	Appliance-based	Downloadable software	SaaS, Appliance	Cloud-based service
RDBMS limitations, data fragmentation	RDBMS limitations, data fragmentation	Full-text index, cost of repeated parsing	Storage backend tradeoffs unknown	Seamless, transparent scalability
One schema, parse at collection, limited SDK	Schema per device, limited device support	Parsing on access via regular expressions	Schema per device	Intelligently evolving log parsing
Limited context model, only network	No context model	No context model	Limited context model, only network	User, network, service model, extensible
No intelligence shared among customers	No intelligence shared among customers	No intelligence shared among customers	No shared intelligence, some SOC service	Global IT log intelligence
Informal community	Informal community	Informal community	Informal Community	Community baked into the service

Go To Market

Self-serve, easy to try, buy and use

Instant gratification

Free trials, freemium model also possible

Value-before-commit

Tiered pricing

Pay for what you use

Web sales, telesales, focused direct touch, channel

Lower cost of sales

Leveraged partnerships

PaaS add-on sale

Economics

Deliver service at high gross margin

Network, Storage, CPU

Inbound network traffic dominates outbound traffic

Storage needs to consider monthly charge due to retention

CPU can be optimized because of elasticity

Tier	Events/Sec	GB/Day	COGS $/Month	MRR*	ACV*		AlertLogic ACV	ArcSight Deal Size
Trial	5	0.5	$5				$2,148	
Silver	120	10	$101	$506	$6,072		$36,000	*$100-500k*
Gold	1,200	100	$1,012	$5,063	$60,756		$153,000	*$0.5-2M*
Platinum	3,000	250	$2,531	$12,656	$151,872		$324,000	*$2-6M*
Diamond	12,000	1000	$10,125	$50,625	$607,500			

*Priced at 80% gross margin

Roadmap

	Release 1 9 – 12 Months	Release 2	Release 3
Intelligence	Anomaly Detection	Pattern Mining	Predictive Analytics
Operations	Troubleshooting	Business Continuity	Service Levels
Security	Threat Analysis	Incident Response	Data Protection
Compliance	PCI Pack	SOX, HIPAA, NERC	Fraud, Risk
Platform	Collection, Search, Reporting, Assets	Workflow, Dashboards, Trending, Identities	Context Modeling, Community

Summary

Cloud-based IT Log Analytics

Large opportunity in growing enterprise market

Team of veteran log management experts

Game-changing *functionality* and ease of adoption