Adetomiwa Oladejo
URN - 6952417

UNIVERSITY OF SURREY

Faculty of Engineering and Physical Sciences Department of Computer Science
MSc programmes in Computer Science

COMM058 – Architectural Thinking for Security

Adetomiwa Oladejo

Adetomiwa Oladejo
URN - 6952417

Table of Contents

Adetomiwa Oladejo

URN - 6952417

Executive Summary

This report sets out a security architecture for Turing Intelligent Energy (TIE), describing how a cloud-based energy optimisation platform can operate securely and reliably in practice. Turing is delivering a service that brings together homeowners, installers, support teams, and external providers within a single operating environment. Protecting sensitive customer and operational data while maintaining continuous availability is a core requirement of this design.

Several architectural challenges shape the design, as the platform marginally relies upon third-party services. Risks around confidentiality and availability are produced as a result. Without clearly defined controls, these characteristics increase exposure to misuse, disruption, or data loss. Addressing these concerns requires deliberate boundary definition, strong identity management, and mechanisms that limit the impact of failures or compromise.

Architecture decisions presented in this report respond directly to these challenges. Structured threat modelling informs security controls, whilst functional and non-functional requirements shape both design and deployment choices. Security, resilience, and accountability are embedded from the beginning, demonstrating that a robust operating model can be achieved without sacrificing usability or practicality.

Architecture Overview

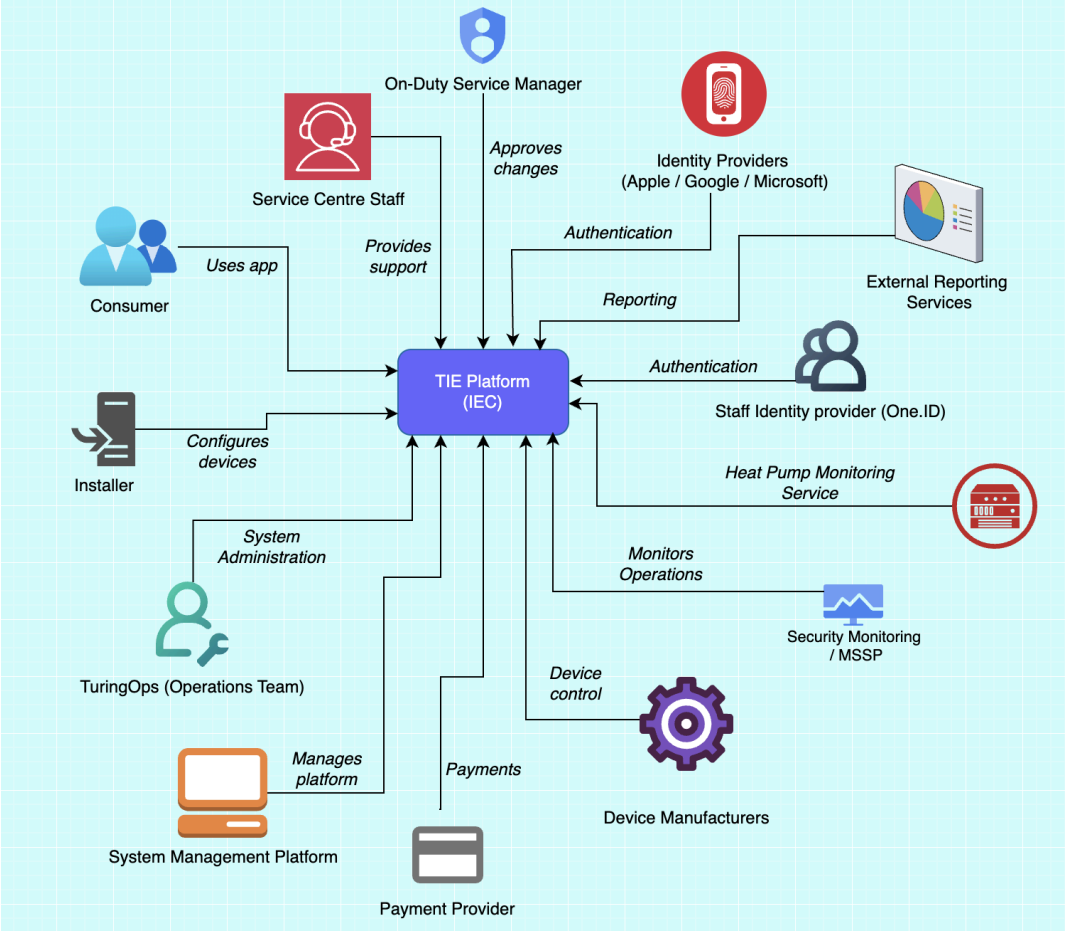

Figure 2: Architecture Overview Diagram

Adetomiwa Oladejo

URN - 6952417

Requirements

Requirements lay out the premise of expectations for the TIE system. It encapsulates factors such as the security expectations, boundaries for the system and other aspects of the scope that must be met. Context, data, and process views are brought together to provide an understanding of how the system works, along with what is protected.

System Context Diagram

A system context diagram provides an outline of a system's boundary and the external actors it interacts with. Through a focus on relationships and trust boundaries as opposed to internal details, a system context diagram aids stakeholders in understanding how the system fits within its wider operational and security environment.

Adetomiwa Oladejo
URN - 6952417

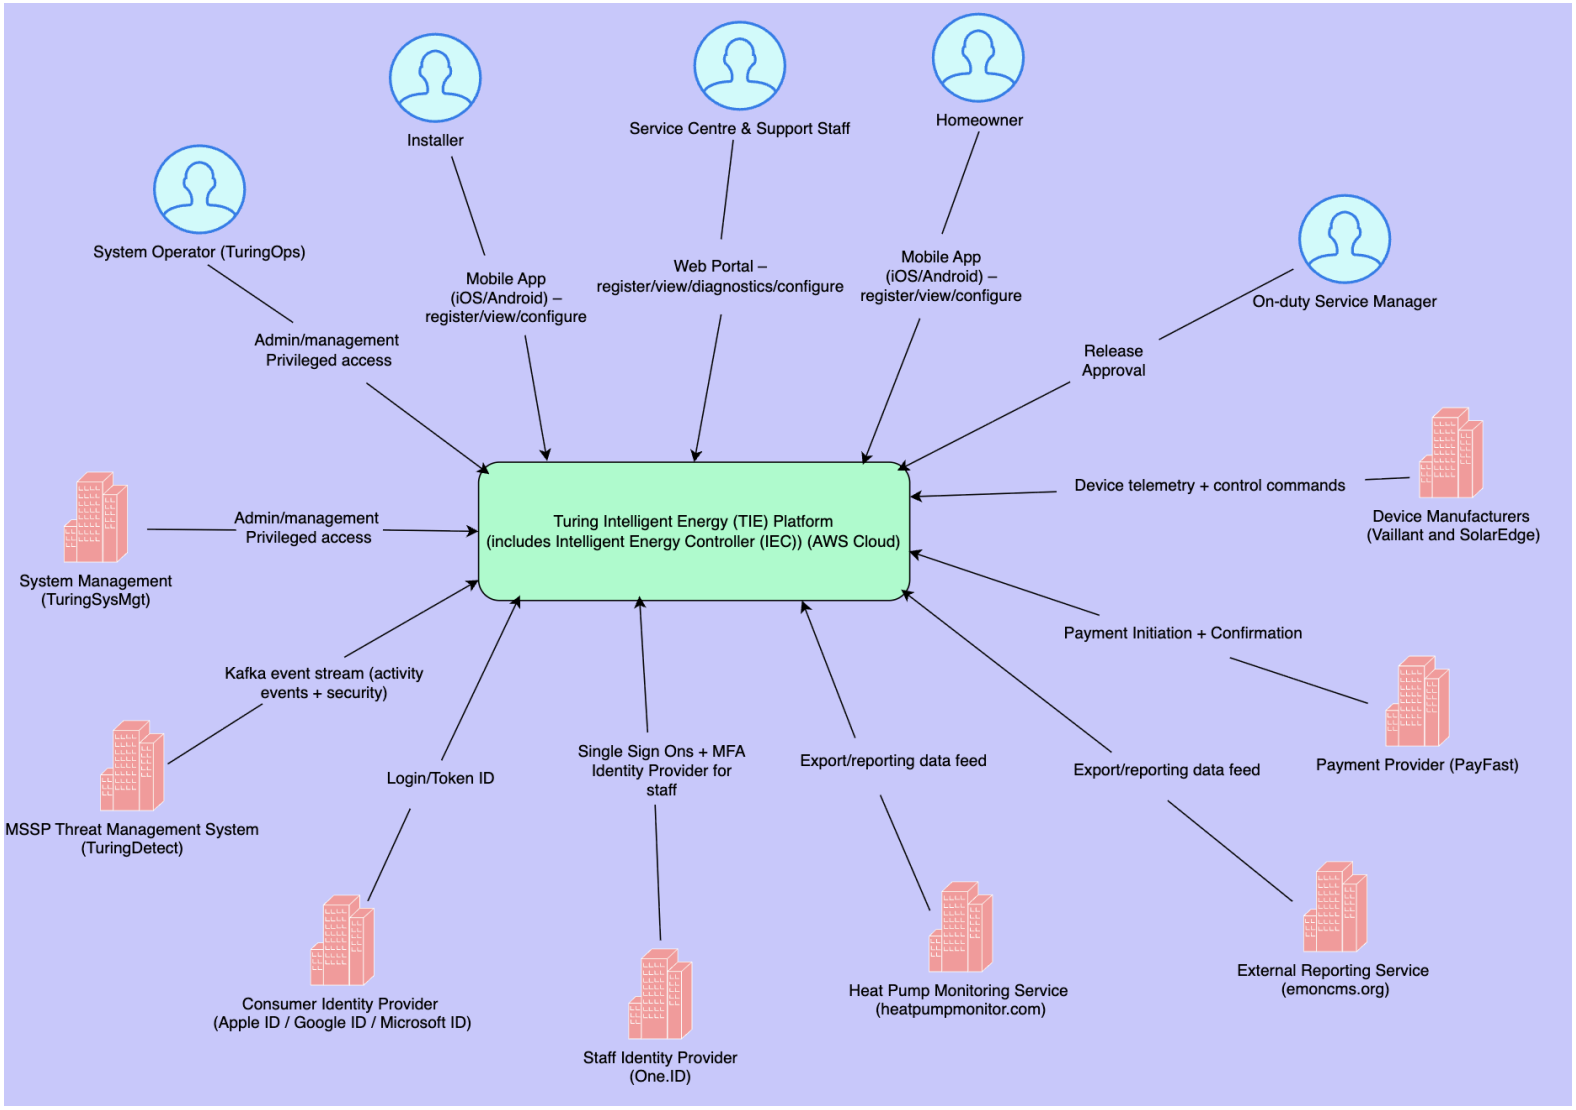

Figure 2: System Context Diagram

Adetomiwa Oladejo

URN - 6952417

To develop the diagram, a clear boundary around the platform as the core trust domain was first constructed. Following the definition, all users and third-party services were divided into user and system actors. Using this method, a zero-trust policy can be applied, allowing access to be managed by identity and controlled interfaces. As the energy data store operates internally, the decision was made for it to be excluded from the diagram (Buckwell, Van Daele and Horst, 2024).

Actor/Use case/Interference

Actors	Description	Interface
Homeowner	Registers energy devices, views energy usage, and configures system preferences for their home	Mobile Application (iOS / Android)
Installer	Installs and configures energy devices for homeowners	Mobile Application (iOS / Android)
Service Centre & Support Staff	Provides operational support to homeowners, including device registration, diagnostics, and configuration changes required to resolve issues	Web Portal
On-duty Service Manager	Provides governance oversight by approving releases and changes	Web Portal
System Operator (TuringOps)	Operates and maintains the TIE platform, performing administrative and operational tasks using privileged access	Operations Management Interface / TuringSysMgt
Device Manufacturers (Vaillant, SolarEdge)	Systems that exchange device telemetry data and receive control commands from the TIE platform	API Integration
Payment Provider (PayFast)	External payment service responsible for processing	Payment API

	homeowner payments and returning payment confirmations.	
External Reporting Services (emoncms.org, heatpumpmonitor.org)	Recieves exported energy and performance data for reporting and benchmarking purposes	Data Export
User Identity Providers (Apple ID, Google ID, Microsoft ID)	Authenticates homeowners and installers and provides identity tokens used by the TIE platform	Identity provider
Staff Identity Provider (One.ID)	Provides authentication and single sign-on services for internal staff accessing the platform	Single sign-on and identity provider
System Management Platform (TuringSysMgt)	Management system used by TuringOps to administer and manage the TIE platform from a separate cloud environment	Management Interface
MSSP Threat Management System (TuringDetect)	An external security monitoring service which receives security events, supporting both threat detection and incident response	Security Event Stream

Table 1: System Context Table

The system context diagram identifies actors, trust boundaries, and data flows that feed directly into the functional and non-functional requirements, shaping how security controls are applied across the platform.

Adetomiwa Oladejo

URN - 6952417

Data Classification (Confidentiality only)

Data classification will act as a guideline for processed information by the TIE platform. Data has been grouped according to its sensitivity, providing a comprehensive way to understand the confidentiality of different data types and the potential impact if that information were to be accessed or disclosed without authorisation (Buckwell, Van Daele and Horst, 2024).

Data Classification	Description
Public	Information intended for unrestricted access where disclosure does not pose a confidentiality risk. This includes non-sensitive service descriptions and publicly available documentation.
Personal Information (PI)	Data that can be linked to individual homeowners or installers, including identity details and household-level energy usage. Loss of confidentiality could infringe privacy rights and erode customer trust.
Sensitive Personal Information (SPI)	Authentication and identity-related material, such as credentials, tokens, and unique identifiers, is issued by identity providers. Compromisation of this data would directly undermine access controls and system confidentiality.
Financial Information	Payment-related records are exchanged with external payment services, including transaction references and billing confirmations. Unauthorised disclosure could expose customers to financial risk or fraudulent activity.
Confidential Operational Data	Information required to operate and support the platform, such as configuration data, diagnostics, and control commands. Exposure could reveal internal system behaviour or assist targeted attacks.
Security Sensitive Data	Data that can expose operational functions, inclusive of audit records, monitoring data and security logs. Loss of confidentiality could increase the attack surface.

Adetomiwa Oladejo

URN - 6952417

Information Asset Inventory

The information asset register builds on the data classification table by identifying the key information assets processed by the platform and assigning each to an appropriate confidentiality category. This provides a perspective of the types of data handled by the system and the potential impact of their disclosure, supporting informed security design and architectural decision-making.

Data Type	Data Fields/Asset	Data Classification	Legal and Regulatory
Energy Optimisation Outputs	IEC optimisation decisions	Confidential Operational Data	
Staff Identity Data	Staff identities, roles, and SSO attributes (One.ID)	Sensitive Personal Information	PI, GDPR
Device Telemetry	Solar and heat pump	Confidential Operational Data	
User Profile Data	Names and contact details of homeowners and installers	Personal Information	PI, GDPR
Energy Usage Data	Household energy metrics	Personal Information	PI, GDPR
User Identity Data	User IDs and external identity providers	Sensitive Personal Information	PI, GDPR
Audit & Activity Logs	User actions, admin changes, system events	Security Sensitive Data	GDPR
Authentication Data	Session tokens and OAuth tokens	Sensitive Personal Information	PI, GDPR
Device Control Commands	Commands sent to household energy devices	Confidential Operational Data	
Payment Transaction References	Payment confirmations and	Confidential (PCI-related)	

	transaction identifiers		

Table 2: Information Asset Inventory

These classified assets form the basis for threat identification in the threat–risk register, thereby aiding assurance that risks are assessed against data that would have a material impact if compromised.

Functional requirements

Functional requirements are integral to the definition of security-relevant aspects, aiding the description of the system's necessary capabilities in order to protect various assets during systematic operations. Each requirement has been aligned with the CSA Cloud Controls Matrix (CCM) to ensure the security design follows industry-standard cloud security best practices (Cloud Security Alliance, 2021).

ID	Requirement	CSA CCM Controls
FR-01	Role-based access control restricts user actions based on assigned roles	IAM-04, IAM-06
FR-02	Energy usage data is securely stored in the Energy Data Store	DSP-01, DSP-02
FR-03	Security event streams are forwarded to the managed security service provider	SEF-02
FR-04	API requests are sanitised and validated before processing	API-02
FR-05	Users authenticate using external consumer identity providers	IAM-02
FR-06	Secure payment initiation and confirmation are supported between the platform and the PayFast provider	DSP-04

FR-07	Authentication and access events are recorded for audit purposes	LOG-02, LOG-04
FR-08	Service centre staff authenticate via One.ID with mandatory multi-factor authentication	IAM-03
FR-09	Application and security events are streamed to Kafka for threat detection and monitoring	LOG-03, SEF-01

Table 3: Functional Requirements

Non-functional requirements or Architectural Characteristics

The purpose of non-functional requirements is to outline the quality characteristics and operating constraints. Expected levels for areas such as performance, availability and operational support are included to guarantee the platform remains both reliable and secure.

ID	Requirement
NFR-01	The platform is designed for 24×7 operation, with operational support provided according to defined service hours and escalation arrangements
NFR-02	Confidentiality of all data must be maintained through controlled access and encryption
NFR-03	Recovery objectives are defined so that critical services can be restored within agreed timeframes following failure
NFR-04	Energy and operational data must be backed up on a regular basis to prevent loss and support recovery when required
NFR-05	User interactions and API requests should respond within acceptable timeframes

NFR-06	Core TIE platform services are designed to operate continuously, with a target availability of 99.9% to support uninterrupted monitoring and automated energy optimisation
NFR-07	Changes to configuration, integrations, and security controls should be deployable with minimal service disruption

Table 4: Non-functional Requirements

Swimlane Diagram with activity description and separation of duties matrix

Within the swimlane diagram below, the lifecycle of a device configuration change from request through to completion has been traced, illustrating how actions are split across different roles and systems. Emphasis is also placed on how checks, controls, and approvals are built into the process, aiding the maintenance of accountability and preventing misuse.

Adetomiwa Oladejo
URN - 6952417

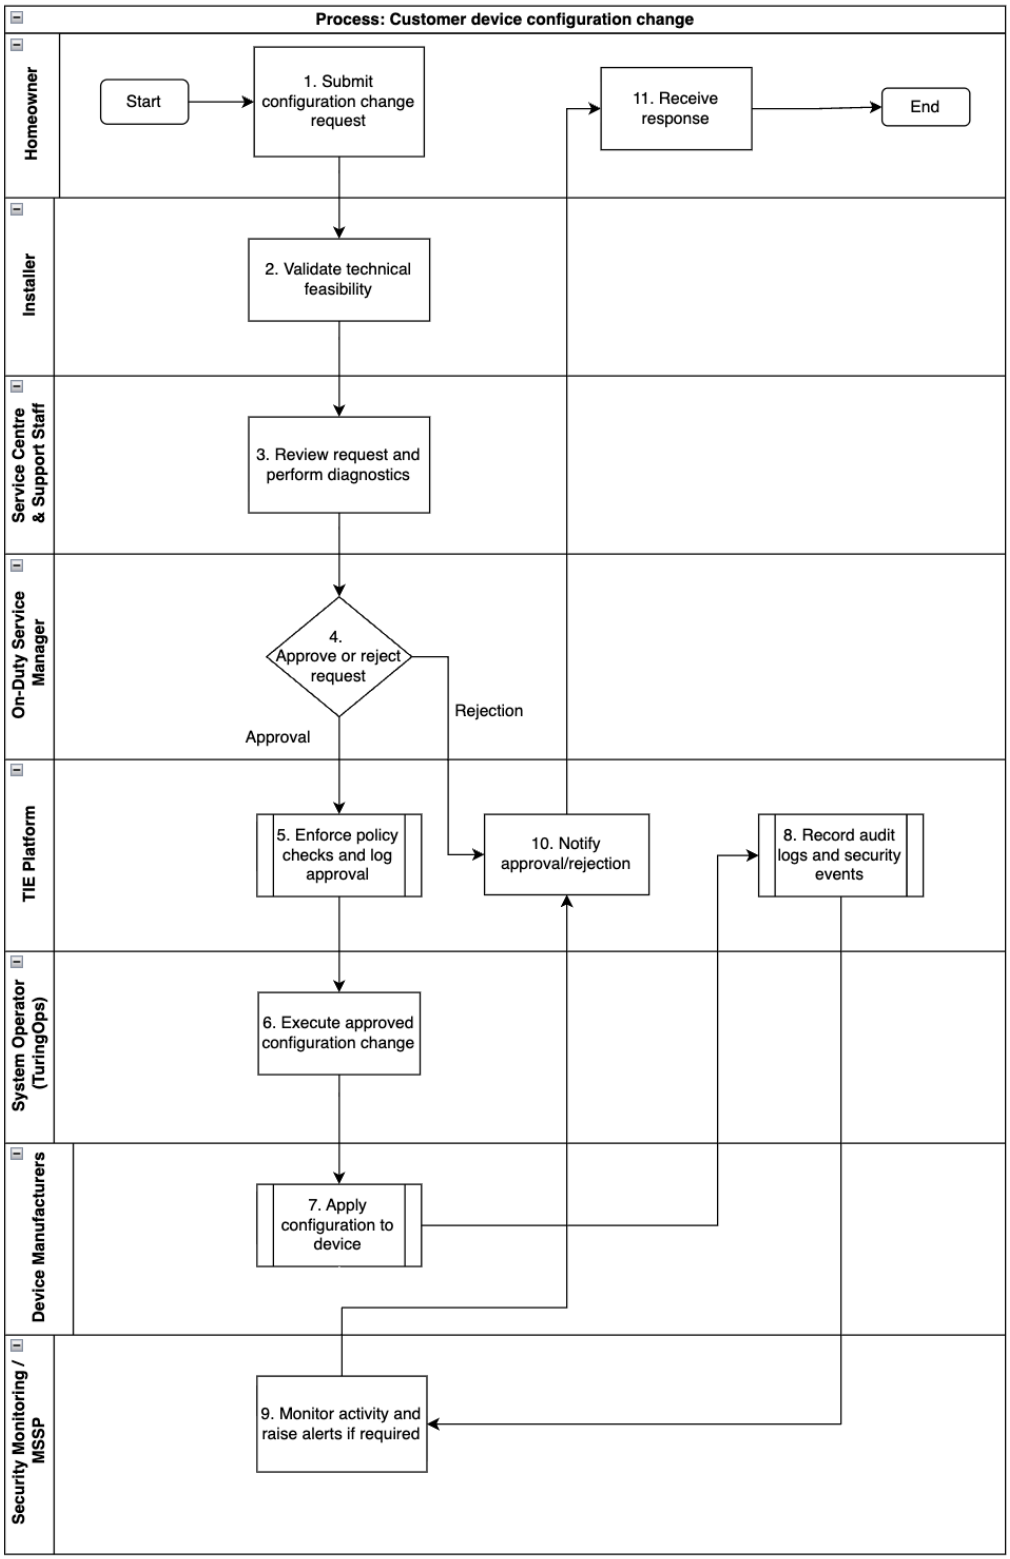

Figure 3: Swimlane Diagram

Adetomiwa Oladejo

URN - 6952417

Separation of Duties Matrix

A separation of duties (SoD) matrix defines which process steps within the customer device configuration change workflow may or may not be performed by the same role. The matrix is utilised to reduce the risk of error, misuse or privilege abuse through assurance of critical actions being appropriately separated across different roles and systems.

SoD Combination	
X	Elevated Risk
*	Low Risk
✓	Combination Allowed

Role	
1	Homeowner
2	Installer
3	Service Centre Support Staff
4	On-Duty Service Manager
5	TIE Platform
6	System Operator (TuringOps)
7	Device Manufacturers
8	Security Monitoring / MSSP

Adetomiwa Oladejo

URN - 6952417

Process Step	Role	ID	1	2	3	4	5	6	7	8	9	10	11
Submit configuration change request	1	1		X		X		X	X				
Validate technical feasibility	2	2	X		*		✓						
Review request and perform diagnostics	3	3		*		*							
Approve or reject request	4	4	X		*			X	X				
Enforce policy checks and log approval	5	5		✓						*		✓	
Execute approved configuration change	6	6	X			X				X	X		
Apply configuration to device	7	7	X			X				X			
Record audit logs and security events	5	8					*	X	X		*		
Monitor activity and raise alerts if required	8	9						X		*			
Notify approval/rejection	5	10					✓						✓
Receive response	1	11										✓	

Table 5: SoD Matrix

Adetomiwa Oladejo

URN - 6952417

Architecture

Architectural design plays a central role in how the Turing Intelligent Energy (TIE) platform operates securely in practice. The focus is on clarifying various aspects of the architecture to display the platform's security and data flows.

Component Architecture Diagram with assets, actors, threat and controls overlay.

Presented within the component architecture diagram are the main building blocks of the TIE platform and the trust boundaries that separate them. By showing where critical services, sensitive data, and security controls lie, risk can be identified so an effective security design can be put in place.

Adetomiwa Oladejo
URN - 6952417

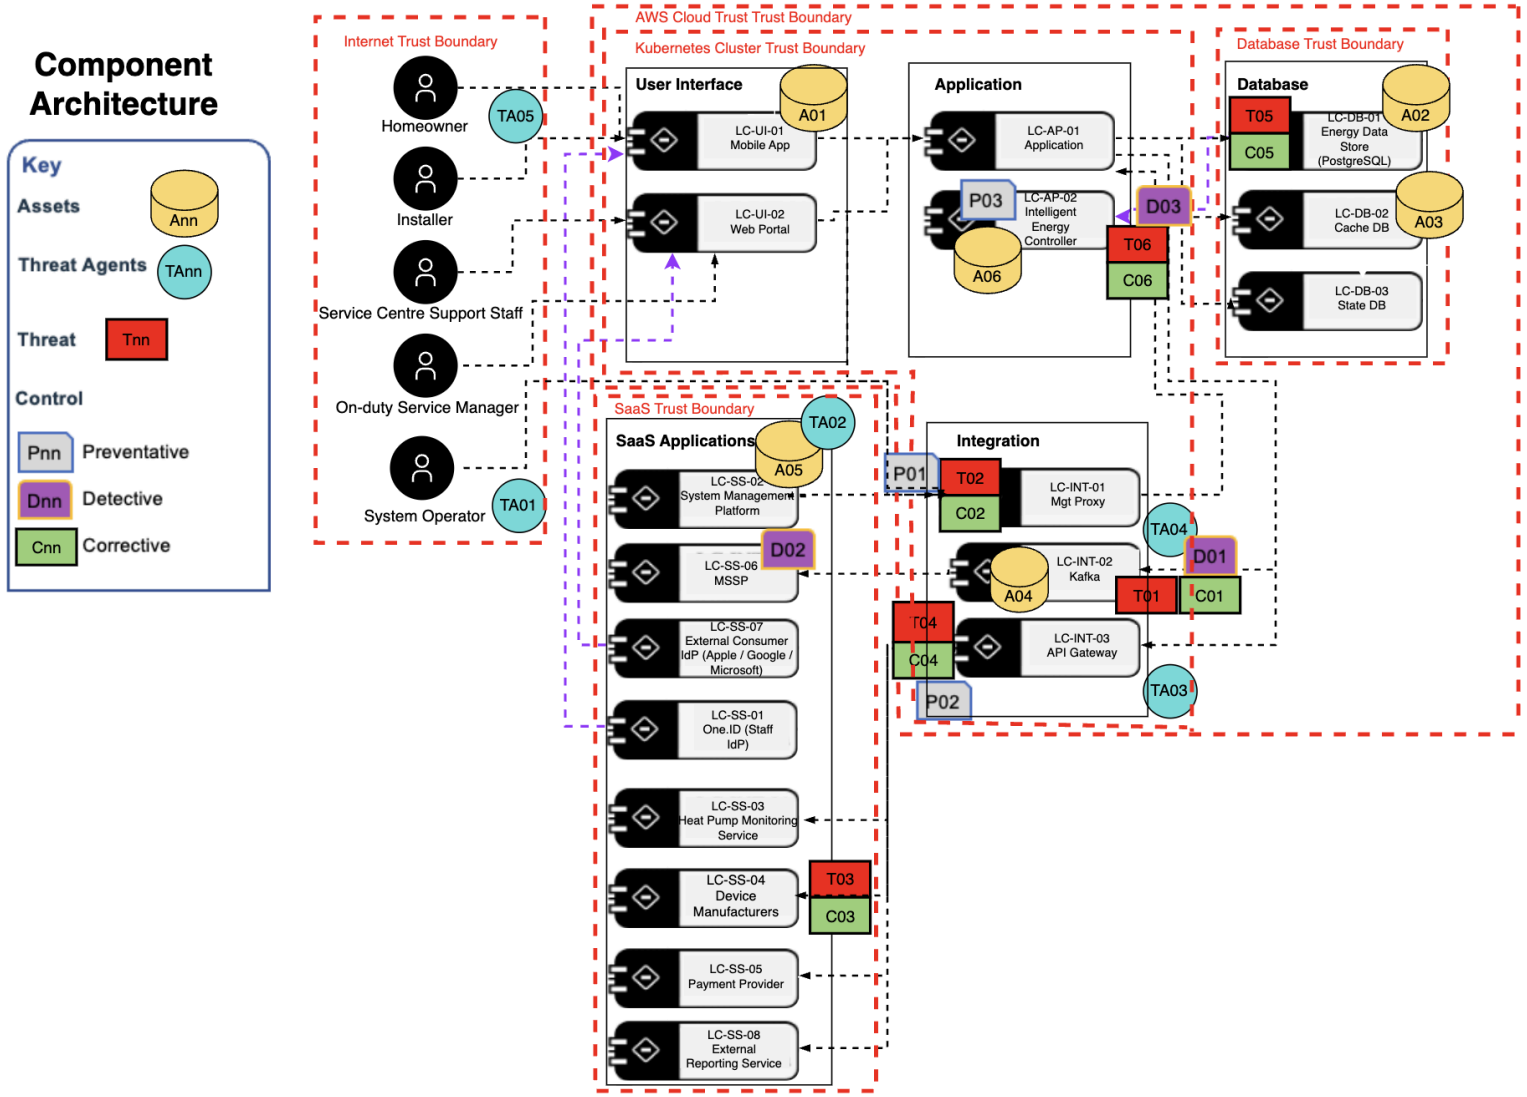

Figure 4: Component Architecture Diagram

Adetomiwa Oladejo
URN - 6952417

ASSETS	THREAT AGENT	THREAT
A01 – UI App Data & Device Configuration A02 – Energy Data Store (EDS) Data A03 – Cache Database Content A04 – Kafka Event Stream Data. A05 – Identity & Privileged-Access Management Artefacts (SaaS boundary) A06 – IEC Decision Logic / AI Model Artefacts	TA01 – Privileged Operator / Outsourcer (TuringOps) TA02 – Compromised SaaS Provider TA03 – External API Attacker TA04 – Adversary Targeting Event / Telemetry Pipelines TA05 – Untrusted End User / Internet Attacker	T01 – Kafka Stream Disruption or Event Injection T02 – Abuse of Management Proxy / Management Channel T03 – Compromise via Device Manufacturer API Integration T04 – API Gateway Attack (Untrusted Invocation / Enumeration / Abuse) T05 – EDS Data Exposure or Unauthorised Access T06 – IEC Manipulation
CORRECTIVE	**PREVENTIVE**	**DETECTIVE**
C01 – Kafka Stream Containment & Recovery ● Quarantine topics / block producer keys / rebuild consumer offsets / replay from known-good checkpoints. C02 – Management Access Revocation & Session Kill ● Disable admin accounts, revoke privileged sessions, rotate management credentials, disable mgmt endpoints temporarily. C03 – Manufacturer API Key Rotation + Device Command Freeze	P01 – Privileged Access Governance via Management Proxy ● Approval workflow enforcement (change record / break-glass), MFA for privileged flows, session controls. P02 – API Gateway Policy & Abuse Prevention ● Strong authN/authZ, schema validation, throttling, request signing where feasible. P03 – IEC Data Integrity Guardrails ● Input validation from EDS to IEC, model governance,	D01 – Kafka Security/Event Logging with Integrity Monitoring ● Detects stream anomalies (unexpected producers, abnormal topic volumes, invalid schemas). D02 – MSSP SIEM Monitoring (TuringDetect) ● Correlation, alerting, incident escalation based on Kafka-fed security events. D03 – IEC + Data Store Audit/Analytics Monitoring ● Detects suspicious optimisation outputs, abnormal query

• Rotate service keys, revoke tokens, disable outbound control commands until trust is restored. C04 – API Gateway Block/Shield Response • Emergency WAF rules, IP/ASN blocks, route to "deny-by-default" policy set, tighten rate limits. C05 – EDS Data Protection Recovery • Revoke DB creds, force password rotation, restore from backup, isolate DB tier, validate integrity C06 – IEC Rollback / Safe-Mode Optimisation • Roll back model/version, switch to deterministic safe ruleset, freeze tuning instructions until validation passes.	restricted write paths for tuning instructions.	patterns, unexpected state changes.

As the component architecture was being designed, the core platform responsibilities were first outlined to gain an understanding of where the trust boundaries are positioned. To aid the creation of a secure-by-design system, the placement of key assets, threats, and controls has been mapped out. Furthermore, for the purpose of enforcing least-privilege access and network segmentation, a separation of the integration, applications, and user access has been applied (Buckwell, Van Daele and Horst, 2024).

Threat-Risk Register with OWASP risk evaluation

A threat–risk register can be used to identify key security risks arising from user activity after inspecting the component architecture. The risks within the register below were assessed using an OWASP-style likelihood and impact approach, with mitigations grouped into preventative, detective, and corrective controls (OWASP, n.d.). Research was conducted to find relevant threats and mitigations for Turing Intelligent Energy (NIST, 2020)(MITRE, 2024).

Adetomiwa Oladejo
URN - 6952417

Table 6: Threat Risk Register

Threat Target	Attack Technique	Threat Actor	STRIDE	Inherent Risk			Risk Mitigation			Residual Risk		
				Likelihood	Impact	Overall Risk	Preventive	Detective	Corrective	Likelihood	Impact	Overall Risk
Integration Layer (API Gateway)	Message poisoning of optimisation event stream	External attacker	DoS	M	H	H	Message sanitisation	Dead letter queue monitoring	Reset users	L	M	L
Intelligent Energy Controller	Abuse of IEC to exhaust system resources	Automated attacker	DoS	M	M	M	Execution limits	Utilisation alerts for resources	Job suspension	L	M	L
Energy Data Store	Energy usage inference via repeated queries	External attacker	Information Disclosure	M	H	H	Query limiting	Query pattern analysis	Restrict queries	M	M	M
User Interface	Credential stuffing	External attacker	Spoofing	H	M	H	MFA	Login anomaly monitoring	Session invalidation	M	L	L
System Management Platform	Privilege misuse via approved or abused admin access	Outsourcer/ Insider	Elevation of Privilege	M	H	M	Break-glass approval	Privileged activity monitoring	Revoke access	L	M	L

Adetomiwa Oladejo

URN - 6952417

The mitigations identified in the threat–risk register directly influence both the architectural controls applied in the component and deployment diagrams.

Deployment Architecture Diagram

From a deployment perspective, the TIE platform operates in a public cloud using segmented VPCs to reduce the attack surface and enforce least-privilege access. It is explicitly shown where trust boundaries sit, along with data flows and the implementation of security controls. The CIA triad is constructed through the application of these factors.

Adetomiwa Oladejo

URN - 6952417

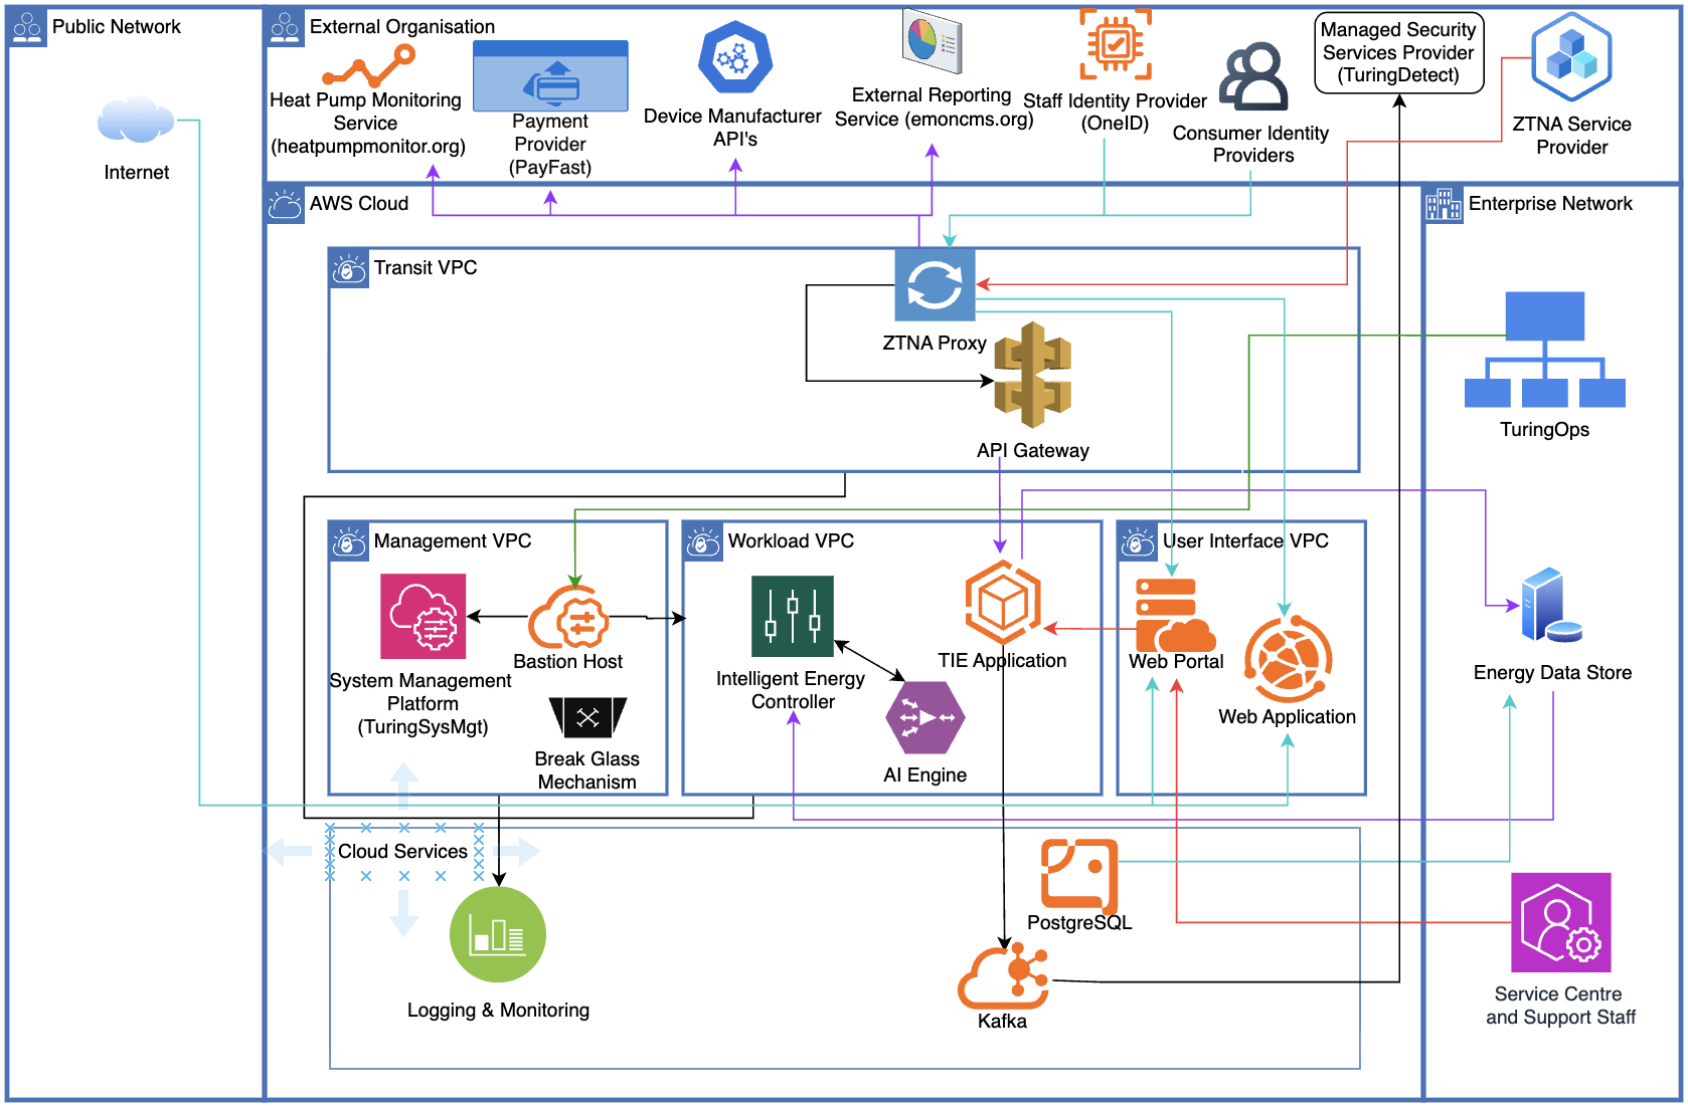

Figure 4: Deployment Architecture Diagram

Adetomiwa Oladejo

URN - 6952417

As the deployment diagram was being illustrated, the idea of security and operational risk was kept in mind. Additionally, to emphasise trust, a Zero Trust Network Proxy (ZTNA) was implemented. The lines have been coloured to create an easier perception of the network and data flows.

Governance Appendix

The role of the Governance Appendix is to assemble the practices that sit behind the technical design. Within this appendix, there are governance mechanisms reported with the objective of supporting the secure operation of the deployed architecture. Architectural controls are reinforced through defined responsibilities, processes, and potential oversight.

Viability Assessment (RAID log)

Viability Assessments outline the main risks, assumptions, issues, and dependencies that may affect delivery and operations. Areas of uncertainty have been made explicit, with clear actions and ownership defined to ensure active management.

Risks

Statement	Action(s)	Owner & Date
There is a risk that reliance on external service providers could impact platform availability, caused as a result of third-party service outages or degraded performance, which may result in loss of service to customers and installers.	Define SLAs, monitor service health, and document fallback procedures.	Owner: Platform Operations Lead Date: Service onboarding milestone
There is a risk that unauthorised access to customer or operational data could occur, caused as a result of misconfigured access controls across multiple cloud trust boundaries, which may result in loss of confidentiality or service disruption.	Enforce identity-based access controls, role separation, and regular access reviews.	Owner: Security Architect Date: Prior to MVP production deployment

Table 7: Risk (RAID)

Adetomiwa Oladejo

URN - 6952417

Assumptions

Statement	Action(s)	Owner & Date
There is an assumption that identity providers used for users and staff authentication remain secure and available. If the assumption is not validated or turns out to be invalid, there is a risk that users are unable to authenticate, which may result in reduced platform availability.	Monitor availability and document incident escalation paths	Owner: Security Lead Date: Before deployment
There is an assumption that cloud-native security services provided by the public cloud platform are correctly configured and maintained. If the assumption is not validated or turns out to be invalid, there is a risk that security vulnerabilities remain undetected, which may result in increased attack surface.	Conduct configuration reviews and security assurance checks	Owner: Cloud Security Engineer Date: Within 4 weeks of the environment build

Table 8: Assumptions (RAID)

Issues

Statement	Action(s)	Owner & Date
There is an issue that the transparency of security events depends on correct event forwarding to monitoring services, which has resulted in delayed detection during testing. This has been caused because event pipelines were not completely validated early in the design.	Validate logging and alerting paths before going live	Owner: Security Operations Lead Date: Within 3 weeks of logging and monitoring integration

There is an issue that some operational tasks currently require elevated privileges, which has resulted in increased reliance on trusted operators. This has been caused because access controls are still being improved.	Introduce tighter role definitions and approval-based privileged access	Owner: Operations Manager Date: Within 1-4 weeks

Table 9: Issues (RAID)

Dependencies

Statement	Action(s)	Owner & Date
There is a dependency on a managed security services provider to monitor security events. If not delivered against in the required timeframe, there is a risk that security incidents will not be detected promptly, which may result in elevated impact.	Define monitoring responsibilities and escalation procedures	Owner: Security Operations Lead Date: SOC onboarding
There is a dependency on external identity providers to authenticate homeowners and staff. If not delivered against in the required timeframe, there is a risk that users cannot access the platform, which may result in service unavailability.	Confirm provider availability and support arrangements	Owner: Security Lead Date: Identity integration milestone

Table 10: Dependencies (RAID)

Adetomiwa Oladejo

URN - 6952417

Architecture Decision Record

This Architectural Decision Record captures a significant security design choice for the Turing Intelligent Energy (TIE) platform. The decision is presented within its architectural and operational context, with consideration given to alternative approaches. Moreover, a clear explanation of why the selected option best addresses the identified risks and design principles is incorporated.

<table>
<tr><td>Subject Area</td><td>Network Security and Access Control</td><td>Topic</td><td>Zero Trust Network Access</td></tr>
<tr><td>Architectural Decision</td><td>Should Zero Trust Network Access be adopted to control access to the TIE platform for staff and privileged users?</td><td>AD ID</td><td>TIE-AD-001</td></tr>
<tr><td>Issue or Problem</td><td colspan="3">Access to the TIE platform is required by users, staff and third-party services across multiple locations and networks. Traditional network-based access models increase exposure by extending implicit trust once connected, which raises the risk of lateral movement, credential misuse, and unauthorised access to sensitive systems.</td></tr>
<tr><td>Assumptions</td><td colspan="3">● Staff, users, and third-party services access the platform remotely over untrusted networks.
● The platform spans multiple trust boundaries within a public cloud environment.
● Access to administrative and operational functions will be tightly controlled and auditable.</td></tr>
<tr><td>Motivation</td><td colspan="3">Controlling who can access which services is a critical aspect of protecting availability and sensitive data. A zero-trust approach reduces the attack surface, using access decisions based on policy, context, and identity.</td></tr>
<tr><td rowspan="2">Alternatives</td><td rowspan="2">Alternative 1: Zero Trust Network Access (ZTNA)</td><td colspan="2">Advantages:
● Restricts access per-application using identity, device, and context
● Reduces attack surface by removing blind network trust</td></tr>
<tr><td colspan="2">Disadvantages:
● Requires changes to access patterns and user workflow</td></tr>
</table>

		• Introduces dependency on identity and policy enforcement services
		Expected effort/Cost • Medium implementation effort with ongoing service costs, offset by reduced security and operational risk
	Alternative 2: Traditional VPN-Based Access	Advantages: • Provides encrypted network access for internal users
		Disadvantages: • Increased risk of lateral movement due to broad network access once granted • Poor compatibility with cloud-native and external services
		Expected effort/cost • Moderate setup cost, but continued costs for monitoring and access management
	Alternative 3: Direct Public Exposure with Identity Controls Only	Advantages: • Simple to implement with the use of existing identity providers and application controls. • Low upfront cost and minimal infrastructure complexity

		Disadvantages: ● Exposes services directly to the internet, increasing the attack surface. ● Limited protection against network-level threats
		Expected effort/Cost ● Low risk initially; however, the risk increases over time
Decision	Zero Trust Network Access will be implemented for staff and privileged access to the TIE platform.	
Justification	ZTNA aligns with the platform's cloud-native design and reliance on remote operational access. The proxy supports least-privilege access and reduces the risk of lateral movement, in conjunction with integrating effectively with identity providers. In contrast to a VPN-based approach, ZTNA offers stronger security with lower risk.	
Implications	● Remote access will be managed through a ZTNA provider. ● Operational processes must incorporate identity and policy-based access controls.	
Derived requirements	● Identity-based access policies for users, staff and external services. ● Centralised logging and monitoring of access decisions.	
Related Decisions	● Use of external identity providers for staff and user authentication.	

Table 11: Architecture Decision Record

Adetomiwa Oladejo

URN - 6952417

At its core, this report defines a security architecture for the Turing Intelligent Energy (TIE) platform shaped by real-world governance requirements, access patterns, and operational responsibilities. A consistent link has been made to all areas of the architecture in order to create a viable solution.

Risks affecting confidentiality and availability are addressed through segmented trust boundaries along with identity-driven access management, with the governance reinforcing these implementations.

By centring design decisions around threat modelling and operational constraints, the architecture is able to avoid unnecessary complexity whilst retaining its resilience. The outcome demonstrates that the TIE platform can remain dependable and secure when protection is treated as a design input.

Adetomiwa Oladejo

URN - 6952417

References

Cloud Security Alliance (2024). *Cloud Security Alliance*. [online] Cloud Security Alliance. Available at: https://cloudsecurityalliance.org/research/cloud-controls-matrix/ [Accessed 20 Dec. 2025].

OWASP (n.d.) *OWASP Risk Rating Methodology.* [online] OWASP. Available at: https://owasp.org/www-community/OWASP_Risk_Rating_Methodology [Accessed: 26 Dec. 2025].

NIST (2020) *Zero Trust Architecture (SP 800-207).* National Institute of Standards and Technology. Available at: https://csrc.nist.gov/publications/detail/sp/800-207/final [Accessed: 26 December 2025].

MITRE (2024) *MITRE ATT&CK® Framework.* [online] MITRE. Available at: https://attack.mitre.org [Accessed 27 Dec 2025].

Buckwell, M., Van Daele, S. and Horst, C. (2024). *Security Architecture for Hybrid Cloud.* Sebastopol, CA: O'Reilly Media.