VervIAM

Security Design Advances

Author: Nya Murray August 2026

Table of Contents

1. Identity & Authentication

VervIAM provides the identity and authentication foundation for Verv applications. Identity is treated as the first element of an end-to-end security context extending across **identity, device, network, application workload and data**, rather than as a standalone application login function.

Account Identity

VervIAM establishes and verifies the authenticated identity context of people, applications and devices before access to protected Verv resources. VervIAM supports access by **people, applications and devices**, allowing the identity model to be applied consistently across different types of endpoint and transaction.

For Verv Safe, the authenticated account holder controls their account, private information, transfer relationships and the initiation of document sending and receiving operations.

Verv Safe account holders can securely send documents to, and request documents from, other people without requiring those correspondents to hold a Verv account.

Account Identity Maintenance and Recovery

VervIAM Account Profile provides authenticated account holders with mechanisms to maintain and recover their account identity without requiring disclosure of additional personal identity information to a third-party identity-verification service.

Changes to protected account information require MFA verification. An account holder may change the registered email address, with notification of the identity change sent to both the existing and replacement email addresses.

VervIAM provides an account recovery path using the account holder's previously established identity information after multi-factor authentication.

These controls are implemented, deployed and tested.

The recovery design balances account recoverability with privacy and security: an account holder is not permanently excluded solely through loss of remembered credentials, while recovery does not require VervIAM to collect or disclose additional identity documentation through an external identity-proofing provider.

VervIAM protects the application identity and authentication process. Security of the user's device and registered email account remains outside the VervIAM application security boundary and is treated as part of the wider end-to-end security context.

Authentication

Verv applications use VervIAM authentication before an account holder is permitted to access protected account and application functions.

Verv Safe uses email-based Multi-Factor Authentication (MFA), with a short-lived authentication code. Authentication state is subsequently subject to VervIAM session and token controls rather than establishing a permanently trusted browser or device.

Sensitive identity changes, including changes to account profile information and authentication details, require identity verification. Changes to the registered email identity use verification of the existing account credentials together with MFA verification before the identity change is completed.

Identity Federation

VervIAM was designed to integrate with existing enterprise identity environments rather than require replacement of them. Its customer console provides configuration of identity federation and protected endpoints. On demand PKI keypair rotation provides an incident response capability.

VervIAM can therefore use external or enterprise identities to establish access to protected resources while retaining VervIAM's security controls at the Verv boundary.

The architecture supports people, applications and devices and can accommodate different credential and federation mechanisms according to the target environment. VervIAM has demonstrated Identity Federation and provides configurable VPN capability for enterprise deployments.

The underlying VervIAM capability is customer configurable through a VervIAM enterprise account. For enterprise deployments, consultancy-led configuration is the preferred delivery model because identity and access architectures and integration requirements differ between organisations. Consultancy is therefore an implementation model rather than a technical requirement for operating the VervIAM console.

Authentication Boundary

Successful authentication establishes identity but does **not by itself grant access to protected application workloads or private data**.

VervIAM separately validates the security context presented when access to a protected endpoint is requested. Current token validity, cryptographic validation and authorisation are therefore distinct from the initial act of authenticating the account holder.

Requests that do not satisfy VervIAM's protected access validation are not admitted to the protected workload.

The mechanisms by which authenticated identity is carried forward, validated against current rotated cryptographic state and exchanged for protected endpoint authority are addressed separately under **Session, Token & Access Control**.

Security Assertion

VervIAM establishes and verifies the authenticated identity context of people, applications and devices before access to protected Verv resources. Authentication establishes identity context - subsequent VervIAM controls independently determine whether that authenticated identity is permitted through the protected access boundary.

2. Session, Token & Access Control

VervIAM separates successful authentication from permission to access a protected application workload. Authentication establishes identity; access is subsequently controlled through short-lived token validation, current cryptographic state, endpoint context and protected credential exchange.

Session Establishment

Verv does not rely on a permanently trusted device or long-lived browser credential to establish continuing access.

Following successful authentication, VervIAM establishes a controlled session in which access to protected resources remains subject to token validity and access checks. Session duration is limited by a six-hour browser inactivity timeout, with the complete sign-in and authentication process required at least once every 24 hours.

Protected Access Boundary

A request to a protected Verv endpoint must pass VervIAM validation before it is allowed to reach the protected application workload.

VervIAM checks that authentication has occurred within the required short validation period and verifies that the token presented for access is valid for the current Verv cryptographic state.

A token that fails this validation is rejected at the Verv front door. The request is therefore not forwarded to the protected workload and cannot use that workload to access protected private data.

This implements the VervIAM Software Defined Perimeter principle that identity and access validation occur before access to the protected resource is permitted.

Protected Credential Exchange

Protected endpoint credentials are retained within the Verv server-side processing environment and are made available only after successful identity, token and access validation. The browser is not provided with the downstream credential required to access the protected endpoint.

The protected credential therefore remains within the trusted Verv/AWS processing environment rather than being exposed directly to the browser or public request path.

Protected forwarding additionally requires a recent authenticated access context.

Token Rotation and Replay Resistance

VervIAM cryptographic material is rotated, and protected requests are validated against the current rotation state.

Protected requests are validated against the current cryptographic state. Authentication material associated with an invalid or expired state fails current access validation.

Token lifetime, authentication recency and current key-state validation operate together to reduce the useful lifetime of intercepted or replayed authentication material.

Endpoint and Access Context

VervIAM can associate access with configured endpoints, scopes and attribute information rather than treating authenticated identity as unrestricted authority.

For enterprise federation, endpoint-specific access context can be configured according to the security requirements of the target environment.

The resulting access decision therefore considers both the authenticated identity and the context in which access to a particular protected endpoint is being requested.

Logout and Session Expiry

VervIAM limits the lifetime of an authenticated session rather than relying upon indefinite browser state.

Session controls include automatic logout after inactivity and an absolute maximum session duration. Logout or expiry terminates the current authenticated access context, requiring authentication to be re-established before further protected operations can proceed.

Security Assertion

VervIAM does not treat authentication as sufficient authority to access a protected workload. Every protected request must satisfy current token, cryptographic and access-context validation at the Verv boundary before it is admitted. Protected endpoint credentials remain behind that boundary and are exchanged only after successful validation.

3. Application Data Protection

Verv protects private application data across its lifecycle rather than relying solely on network transport or storage-layer encryption. Protection is applied at the browser, transport, application-processing and persistence boundaries, with plaintext deliberately constrained to the locations in which it is required.

Browser Data Protection

The authenticated account holder views and enters their private application data in clear text within the browser.

Before private data leaves the browser, Verv applies application-level encryption. The resulting encrypted payload is then transmitted over an independently encrypted TLS connection.

Application encryption and transport encryption therefore provide separate protection layers. TLS protects the communication channel; Verv application encryption protects the private data carried through that channel.

Protected Application Processing

Inbound application traffic passes through Verv's layered network and API security controls before reaching the authorised application-processing environment.

Application-encrypted data is decrypted within the authorised server-side processing environment only when processing requires access to its contents. Plaintext exists transiently during that authorised processing and is not persisted in its intermediate form.

Re-encryption Before Persistence

Private data is not passed directly from its transient plaintext processing state into persistent storage.

Verv generates new cryptographic material and immediately re-encrypts the private data before encrypted persistent storage, where it receives an additional platform-managed at-rest encryption layer.

The encryption protecting persisted private data is therefore distinct from both the temporary encryption used to protect data leaving the browser and the TLS encryption protecting network transport.

Plaintext Boundary

Verv deliberately constrains where private application data can exist in plaintext.

For the account holder, plaintext is presented within the authenticated browser when information must be viewed or entered.

Within the server-side architecture, plaintext exists transiently inside the authorised Lambda execution environment where decryption is required to process and re-encrypt the information.

Private application data is not intentionally persisted in plaintext.

This distinction defines the Verv application plaintext boundary without claiming that encrypted information can be processed without ever being decrypted.

Data Protection in Depth

Verv therefore applies multiple independent protection mechanisms to private application data:

Application protection — private data is encrypted before leaving the browser.

Transport protection — the encrypted application payload is transmitted over TLS.

Processing protection — decryption is constrained to the authorised server-side processing operation.

Persistence protection — private data is re-encrypted using newly generated cryptographic material before database storage.

Platform storage protection — the encrypted application data receives the additional at-rest encryption provided by DynamoDB.

No single one of these controls is treated as a substitute for the others.

Security Assertion

Verv protects private application data independently of transport and storage encryption. Private data is encrypted before leaving the browser, transmitted over TLS, decrypted only within the authorised server-side processing boundary when required, re-encrypted using newly generated cryptographic material before persistence, and maintained in encrypted storage. Plaintext is deliberately constrained to the authenticated browser and transient authorised application processing.

4. Cryptographic Key Management

Verv separates cryptographic key functions according to their security purpose and required lifetime. Ephemeral encryption, service-level PKI protection and account-specific encryption are distinct mechanisms rather than uses of a common application key.

Ephemeral Data Protection

Ephemeral cryptographic keys are generated as required to protect private application data moving between the browser and the authorised Verv processing environment.

These keys provide application-level encryption in addition to TLS transport protection. They exist for the immediate data exchange and are discarded after use rather than retained as long-lived application credentials or storage keys.

This separates protection of data in motion from the cryptographic material used to protect Verv services and persistent security configuration.

PKI Service Protection

Verv uses a public/private PKI keypair to protect service endpoints and their associated security configuration.

The PKI layer provides cryptographic protection for sensitive service information required to establish and operate protected Verv connections. Protected service configuration is not exposed directly to the requesting browser and is made available within the authorised server-side environment only as required for validated access.

Account-Specific Configuration Protection

Verv applies account-specific cryptographic protection to sensitive configuration data. This protection is separate from the higher-level cryptographic controls protecting the service environment, maintaining separation between account and service security boundaries.

Service Rotation

For Verv Safe, protected service configuration is automatically rotated daily together with its associated PKI keypair.

Rotation replaces the cryptographic state protecting the service rather than allowing the same service configuration and key material to remain indefinitely valid. Current cryptographic state also participates in the VervIAM protected-access validation described under **Session, Token & Access Control**.

In a customised enterprise VervIAM deployment, the rotation interval can be configured according to the security requirements and risk profile of the protected service. More sensitive or rapidly changing environments can therefore use substantially shorter rotation periods. This configurability is an enterprise VervIAM capability and is not exposed as a configurable option within Verv Safe.

Enterprise VervIAM also supports on-demand PKI keypair rotation, allowing current cryptographic state to be replaced as an incident-response measure when required.

Separation of Cryptographic Functions

The Verv cryptographic architecture therefore separates:

Ephemeral encryption — temporary keys supplement TLS to protect private application data during exchange and are discarded after use.

Service PKI — a public/private keypair protects Verv service endpoints and their security configuration.

Account configuration encryption — account-specific cryptographic protection is applied to sensitive configuration data independently of service-level protection.

Persistence encryption — newly generated cryptographic material protects private application data before database persistence, as described under Application Data Protection.

Platform encryption — managed storage services provide an additional encryption layer for persisted information.

Rotation — Verv Safe service configuration and its associated PKI keypair are automatically replaced daily.

These mechanisms have different purposes and lifetimes and are used together as complementary security controls.

Security Assertion

Verv does not depend upon a single long-lived encryption key across the application architecture. Ephemeral keys protect private data exchanges and are discarded after use; account-specific cryptographic protection separates sensitive account configuration from service-level protection; and PKI protects service endpoints and security configuration. Verv Safe automatically rotates its protected service configuration and associated PKI keypair daily; customised enterprise VervIAM deployments can apply scheduled or on-demand rotation appropriate to the security requirements of the protected service.

5. Document Transfer Security

Verv Safe provides controlled transfer of documents between an account holder and their correspondents without requiring the correspondent to hold a Verv account.

Document upload and retrieval are first subject to Verv/VervIAM access validation. Only after the relevant transfer context has been validated is temporary, scoped storage authority provided for the specific upload or download operation.

Secure Upload and Retrieval

A participant cannot directly obtain general access to the Verv Safe document store.

The relevant Verv access controls must first be satisfied for the Document Transfer. Once the transfer request has been validated, Verv Safe provides temporary, scoped authority for the required upload or download operation over HTTPS.

The temporary authority therefore provides the final, time-limited storage capability within a transfer that has already passed the Verv security boundary.

Persistent storage credentials are not exposed to the browser or transfer participant.

Request-Specific Personal Data Protection

Each Document Transfer generates a new temporary secret associated specifically with that request.

This temporary secret is used only to encrypt and decrypt the personal data associated with that transfer. It is generated afresh for each new request and is not reused as a general account or service credential.

This separates the protection of transfer-related personal data from the mechanism used to transport the document payload itself.

Separation of Access and Document Transport

Verv Safe applies VervIAM application-level protection to the personal data associated with the transfer.

The document payload itself is uploaded and retrieved using temporary, scoped storage authority over HTTPS after the Verv access boundary has been successfully passed.

The current standard Verv Safe implementation does not add a separate Verv application-layer encryption mechanism to the document contents during the transfer.

Document transfer security therefore combines:

Verv access validation — determines whether the participant may proceed with the transfer.

Request-specific encryption — protects the personal data associated with that transfer.

Transfer authority — provides temporary, scoped permission for the required document upload or download.

HTTPS transport protection — protects the document while it is transmitted.

Platform storage protection — protects the document while temporarily stored within the Verv Safe transfer ecosystem.

Transfer and Storage Access

Verv Safe does not expose its underlying document store as a browsable repository.

Transfer participants receive only the temporary authority required to perform the authorised document operation after the Verv security checks for that transfer have succeeded.

Persistent storage credentials and unrestricted storage access remain outside the participant-facing transfer path.

Temporary Document Custody

Verv Safe treats document storage as part of the transfer process rather than as permanent document storage.

Documents are retained only for the defined transfer lifecycle and are subsequently subject to automatic deletion.

The detailed retention, expiry and destruction controls are addressed under **Storage, Retention & Destruction**.

Security Boundary

Verv Safe does not rely upon possession of temporary storage authority as its sole document-transfer security control. Access to the upload or retrieval process is first governed by Verv/VervIAM validation, after which temporary scoped storage authority is provided to complete the authorised transfer.

The standard Verv Safe service does not claim separate application-layer encryption of the document payload itself.

Security of the participant’s own browser or device remains outside the Verv Safe application boundary.

Security Assertion

Verv Safe protects document transfer through layered access control. A participant must first satisfy the relevant Verv/VervIAM transfer validation before temporary, scoped storage authority is provided. Documents are then uploaded or retrieved over HTTPS using that limited authority, while transfer-related personal data is separately protected using a newly generated request-specific secret.

6. Storage, Retention & Destruction

Verv Safe separates temporary document custody from the application information required to identify, manage and audit a Document Transfer.

Documents are stored only for the period required to complete the transfer and are automatically deleted according to the Verv Safe transfer lifecycle. Application and account information is retained separately according to its operational and security purpose.

Temporary Document Storage

Verv Safe is designed as a document transfer service rather than a permanent document repository.

Documents uploaded for sending or receiving are held temporarily within protected storage so that the intended transfer can be completed. Verv Safe document repositories apply automatic lifecycle expiration after the defined two-day retention period. Deletion is managed by the storage lifecycle and may occur asynchronously after a document becomes eligible for deletion.

Access to those documents remains subject to the Verv/VervIAM transfer controls and temporary storage authority described under **Document Transfer Security**.

The underlying document store is not exposed to transfer participants as a general browsable repository.

Automatic Document Deletion

Temporary document storage has a defined lifecycle.

Verv Safe automatically deletes transferred documents after the applicable transfer period rather than relying upon the account holder or recipient to remember to remove them manually.

Automatic deletion reduces the period during which transferred confidential information remains within the Verv Safe storage environment and therefore reduces unnecessary persistence of document data after its transfer purpose has ended.

Separation of Documents and Transfer Records

Deletion of a transferred document does not require deletion of all information identifying that a transfer occurred.

Verv Safe maintains the application information required to identify and manage transfers separately from the temporary document payload. This allows an account holder to view transfer status and

history, identify sent and received transfers, reuse recipient information and copy appropriate information between transfers without requiring the original document to remain stored. **A previous transfer can also be re-initiated from its retained transfer record without requiring the original document payload to remain stored.**

Private application information retained for these purposes remains subject to the VervIAM data-protection controls described under **Application Data Protection**.

Account and Service Lifecycle

Verv applies lifecycle controls not only to individual document transfers but also to application services associated with an account.

Information that is no longer required following expiry or termination of the applicable service is subject to automated cleanup according to the service lifecycle.

This separates operational retention from indefinite retention: information is maintained where it has a continuing account, transfer, security or service purpose and can be removed when that purpose ends.

Security Logs and Audit Information

Security and access records have a different purpose from temporary document storage.

Verv maintains authentication and access information required to provide account-visible access history and support security monitoring and auditability.

Such security records are therefore governed separately from the lifecycle of an individual transferred document.

The security logging and audit model is addressed further under **Privacy, Auditability & Assurance**.

Data Minimisation by Lifecycle

Verv Safe's retention architecture applies the principle that different information has different required lifetimes.

Document payloads — temporary custody for the transfer lifecycle.

Transfer information — retained where required to identify and manage transfer activity.

Reusable recipient information — retained where the account holder elects to reuse transfer relationships.

Security records — retained for security monitoring and audit purposes.

Account and service information — retained according to the account and service lifecycle.

This avoids treating every item associated with a Document Transfer as though it requires the same retention period.

Security Assertion

Verv Safe is designed for temporary document custody rather than permanent document storage. Transferred documents are automatically deleted according to the defined transfer lifecycle, while transfer, account and security information is retained separately according to its continuing operational or audit purpose. This separation reduces unnecessary persistence of confidential document payloads without removing the information required to manage and account for transfer activity.

7. Network, API & Platform Security

Verv applies layered network, API and cloud-platform controls around its application workloads. Public application traffic is delivered through AWS content-delivery and API services over encrypted transport, while application endpoints are subject to access, browser-origin and web-application security controls before reaching the processing layer.

The underlying cloud platform is provided and managed by AWS, with Verv responsible for the configuration and application-level controls deployed on that platform.

Network and API Path

Public Verv service domains are managed through Amazon Route 53 and routed to CloudFront distributions. AWS-managed certificates provide TLS protection for the public service endpoints. CloudFront then routes authorised application traffic through the configured AWS service path to API Gateway and the protected application workloads.

Verv application traffic is delivered through the AWS network and service environment using CloudFront and API Gateway.

Communications use TLS 1.2 or higher to protect network transport between the participant and the Verv service.

API Gateway provides the controlled API ingress path to the serverless application workloads rather than exposing the underlying Lambda processing functions directly as the normal public application interface.

AWS WAF protection provides an additional defensive layer against malicious web requests before traffic reaches protected application processing.

Browser-Origin Control

Verv configures Cross-Origin Resource Sharing (CORS) controls for its browser-accessible APIs.

CORS policies restrict which browser origins are permitted to make supported cross-origin requests to Verv API resources. Enforcement of CORS policy is performed by conforming browser implementations and is therefore treated as a browser-origin control rather than as a substitute for VervIAM authentication, token validation or server-side authorisation.

Verv does not rely upon CORS as an identity or access-control mechanism.

Defence in Depth

Network and API controls operate together with the identity, token and application controls described elsewhere in the Verv security posture.

A request may therefore encounter multiple independent control boundaries, including:

TLS transport protection — protects network communications.

CloudFront — provides the public content-delivery/service boundary.

AWS WAF — provides web-application request filtering.

API Gateway — provides controlled API ingress to application services.

CORS — restricts permitted browser-origin interaction with supported APIs.

VervIAM — performs identity, token, cryptographic and protected-access validation.

Application controls — constrain processing and protection of private data within the authorised workload.

These controls are complementary; successful passage through one layer does not replace validation required by another.

AWS Managed Platform

Verv application workloads operate on AWS managed cloud services.

AWS provides the underlying cloud infrastructure and managed-service security capabilities, while Verv configures and operates the Verv-specific identity, application, access, cryptographic, data and lifecycle controls deployed using those services.

This distinction preserves the cloud shared-responsibility boundary: use of an AWS managed service does not replace Verv's responsibility for secure application design and configuration.

Enterprise Private Network Access

Customised enterprise VervIAM deployments can add configurable private network access where required by the customer's security architecture. This capability can combine enterprise directory controls with VPN access as an additional boundary around the protected Verv environment.

Layered Enterprise Integration

The optional VPN model demonstrates that VervIAM access controls can operate in conjunction with existing enterprise network and directory security rather than replacing them.

An enterprise deployment can therefore combine its own directory and private-access requirements with VervIAM identity, token, endpoint and application controls according to the security requirements of the protected environment.

Email Service Security

VervIAM service email is delivered through Amazon SES. The VervIAM domain is configured with DKIM and DMARC, providing authenticated domain alignment for legitimate VervIAM service email and visibility of attempted unauthorised use of the domain. DMARC reports are delivered automatically and reviewed daily as part of Verv security monitoring.

Security Assertion

Verv protects application access through layered network, API and platform controls rather than relying upon a single perimeter. Standard Verv services use encrypted transport, CloudFront, AWS WAF-protected API Gateway, configured browser-origin controls and VervIAM protected-access validation before requests reach application workloads. AWS provides the managed cloud platform, while Verv retains responsibility for its application and security configuration. Custom enterprise VervIAM deployments can additionally integrate directory-controlled private network access where required.

8. Enterprise Identity & Private Access

VervIAM is designed to integrate with existing enterprise identity and access environments rather than require their replacement.

Enterprise identity federation, protected endpoints, credentials and access context can be configured according to the security architecture of the organisation and the requirements of the protected application environment.

Enterprise Identity Federation

VervIAM can federate with enterprise identity-management systems, including Microsoft Entra ID, allowing an organisation's existing identity environment to participate in access to VervIAM protected resources.

The federation approach can be customised according to enterprise requirements rather than imposing a single identity provider, credential model or integration pattern.

VervIAM has been demonstrated with a Microsoft enterprise identity integration , establishing the practical ability to integrate an external enterprise identity environment with VervIAM protected access.

Federated Credential Protection

VervIAM supports different mechanisms for protecting federated identity and access information according to the integration being implemented.

Federated credentials can be held in encrypted form where protected credentials are required by the target endpoint.

Where federation uses a JWT, the token can instead be cryptographically signed and issued with a short validity period appropriate to the access requirement.

This allows the protection mechanism to reflect the nature of the federated interface rather than treating all enterprise credentials as equivalent.

Configurable Enterprise Access

VervIAM provides a customer-configurable service interface for establishing and managing protected REST services and enterprise endpoints. This allows enterprise access requirements to be configured through VervIAM rather than requiring application-specific changes to the underlying VervIAM security architecture.

For each protected service, authorised customers can configure the identity context, endpoint access requirements, credential mechanism and security attributes required by the target environment. VervIAM supports people, application, device and token-based identity contexts and can accommodate different federation and credential models according to the protected service.

Connection parameters, endpoint identity information and access attributes can be protected and associated with the individual service. Endpoint-specific access requirements can therefore be configured independently rather than requiring every connected enterprise application to adopt an identical identity or access model.

Configured services can subsequently be invoked through VervIAM or integrated into an application while remaining subject to the VervIAM protected-access boundary.

Private Network Access

Where an enterprise requires an additional private network boundary, VervIAM can operate in conjunction with enterprise directory and VPN controls.

VervIAM has been implemented with an AWS Active Directory-controlled VPN model in which access to the Verv federation environment is dependent upon the user being present in the authorised directory and connecting through the required VPN.

In this configuration, network access is logged and the protected Verv environment does not permit the user to sign in unless the required VPN connection is active.

The VPN capability is an optional enterprise VervIAM integration and is not required for, or included as part of, the standard VSafe service.

Customer Configuration

VervIAM can be customer-configured directly through its configuration console.

The VervIAM configuration console allows an authorised customer to create and manage protected REST services, configure endpoint credential and encryption requirements, associate identity information and endpoint attributes, and review service access logs.

Configured services can be invoked through the VervIAM portal or integrated into an application.

VSafe uses this underlying VervIAM security architecture but does not expose the enterprise configuration console as part of the standard VSafe application.

Consultancy-Led Integration

Although VervIAM provides configurable enterprise identity and endpoint capability, enterprise federation frequently involves organisation-specific identity architecture, claims, credentials, network controls, application interfaces and access requirements.

VervIAM can therefore also be delivered through a consultancy-led engagement in which the federation and protected-access model are configured and integrated according to the customer's particular environment.

This approach allows VervIAM to extend and protect existing enterprise identity infrastructure rather than requiring a rip-and-replace identity programme.

Reference Architecture

The VervIAM **Cloud Transformation Security Reference Architecture** provides the broader architectural context for integrating identity, infrastructure, application and data security across cloud and enterprise environments.

The reference architecture treats IAM as an end-to-end security function spanning device, network, application and data controls and supports integration with existing identity and access services.

Security Assertion

VervIAM is designed to federate with and extend existing enterprise identity environments rather than replace them. Enterprise identity and protected-access controls can be configured according to the target environment, with credential protection and short-lived federation mechanisms applied as appropriate. Optional private network controls can add directory-controlled VPN access, while direct customer configuration and consultancy-led integration support organisation-specific security requirements.

9. Threat & Abuse Resistance

Verv applies multiple preventive and detective controls to reduce the ability of a malicious or unauthorised request to progress from the public network boundary to a protected application workload or private data.

Threat resistance is layered across network ingress, authentication, token validation, cryptographic state, endpoint authorisation and monitoring rather than relying upon any single control to establish trust.

Preventive controls are supplemented by scheduled security assessment and alerting to identify unauthorised or anomalous access activity.

Layered Admission Control

Before protected forwarding is permitted, VervIAM validates the presented token, the current cryptographic state associated with the protected service, and the recency of the authenticated forwarding context.

Token validity alone is not sufficient. VervIAM additionally requires the most recent relevant authenticated forwarding context for the account to fall within a 10-second validation window before protected forwarding is permitted.

A request that fails any required validation is not forwarded to the protected application endpoint.

Token Replay Resistance

VervIAM combines conventional token-expiry controls with an independent authentication-recency check. Even where presented token material remains within its nominal validity period, protected forwarding requires the relevant authenticated forwarding context to satisfy the separate 10-second validation window and the token to validate against the current protected-service cryptographic state.

These independent temporal and cryptographic controls reduce the useful opportunity for replay of previously obtained access material.

Endpoint-Specific Authorisation

VervIAM can associate protected services with endpoint-specific identity context, scopes, tags and credentials.

Authentication therefore does not imply universal authority across configured services. Access can be constrained according to the identity and access information associated with the particular protected endpoint.

Where federation uses a JWT, signed short-lived federation tokens can be configured according to the requirements of the target service.

Sensitive endpoint configuration, scopes and tags can be stored in encrypted form rather than exposed through the public application path.

Network and Application Threat Resistance

Requests reaching Verv services are subject to the network and API protections described under **Network, API & Platform Security**.

AWS WAF provides web-application request filtering, while AWS edge and platform protections provide additional resistance to common network and denial-of-service attacks. Verv's public security

architecture also incorporates network and infrastructure monitoring to provide visibility of allowed, denied and potentially malicious activity. (Verv IAM)

These controls provide defence in depth before a request reaches the Verv application workload.

Credential and Secret Exposure Reduction

Verv reduces the exposure of persistent credentials by keeping protected endpoint credentials and service configuration behind the Verv security boundary.

Temporary cryptographic material is used where temporary protection is sufficient, while persistent endpoint authority is not distributed to browsers merely because a user has authenticated.

For document transfers, temporary presigned storage authority is issued only within the validated Verv transfer process rather than exposing persistent AWS storage credentials.

Reducing the number, exposure and useful lifetime of credentials limits the value of material that could otherwise be intercepted or reused.

Security Monitoring and Detective Controls

Verv supplements preventive security controls with platform-wide detective monitoring across the AWS Verv environment.

Custom CloudWatch dashboards provide visibility at both the VervIAM security-control and application-service levels. Authentication monitoring includes token-validation outcomes, protected endpoint access, authentication-access logging, validation latency and failed authentication requests. Application monitoring includes API errors, request and integration latency, service activity and tracing.

Relevant CloudWatch alerts are configured to identify conditions requiring investigation, while VervIAM authentication and access records provide an additional application-level audit trail.

AWS security monitoring and assessment services provide further platform-level detective controls. Together these mechanisms allow security activity to be observed across the Verv environment rather than only within an individual application such as VSafe.

Trust Boundary

Verv protects the network, identity, application and data paths within the Verv-controlled service boundary.

It does not claim that application controls can make a compromised participant device trustworthy. Security of the participant's operating system, browser environment and local endpoint remains outside the Verv application boundary.

The architecture instead limits what a successful interaction from that endpoint can reach and how long temporary access authority remains useful.

Compound Access Resistance

Verv protected services do not expose a general root-level application API from which protected resources can be enumerated. Requests are directed to resource-specific service paths and must provide the correctly structured parameters required by the target operation.

Knowledge of a service endpoint is not itself sufficient authority to use it. A protected request must additionally satisfy VervIAM token, current cryptographic-state and recent authentication-context validation before forwarding is permitted.

These controls are deliberately cumulative. Failure of any required validation prevents the request from progressing to the protected workload.

Security Assertion

Verv resists unauthorised and abusive access through layered controls spanning network ingress, token validity, independent 10-second authentication-recency validation, current cryptographic state, endpoint-specific authorisation, credential isolation and security monitoring. VervIAM does not equate an unexpired token with sufficiently recent authentication for protected forwarding.

Authentication alone does not establish authority to reach a protected workload, and failed access validation prevents forwarding to the protected endpoint. Preventive controls are supplemented by logging and monitoring to provide visibility of attempted and successful access.

Verv's logging and monitoring architecture is designed for customisation. Deployment-specific dashboards and alerts can be created from the available application, authentication and AWS platform telemetry, and the logging architecture can be integrated with an enterprise's existing security-monitoring environment where required.

10. Privacy, Auditability & Assurance

VervIAM treats privacy, auditability and security assurance as related but distinct requirements.

Private information is limited to what is required to provide the service; security-relevant activity is recorded so that access can be understood and investigated; and preventive controls are supplemented by continuing monitoring and assurance processes designed to establish what is actually occurring within the Verv environment.

Privacy by Architectural Constraint

Verv seeks to minimise unnecessary collection, exposure and persistence of private information through the architecture of the service.

Private application data is protected according to the controls described under **Application Data Protection**, while document payloads are retained only temporarily according to the lifecycle described under **Storage, Retention & Destruction**.

VSafe correspondents do not need to become Verv account holders merely to participate securely in a Document Transfer.

The architecture therefore avoids making collection of additional identity or behavioural information a prerequisite where it is not required to perform the service.

User Visibility and Accountability

Verv provides account holders with visibility of security and application activity associated with their account.

Authentication and access logs can be viewed through the Verv Account Profile, allowing the account holder to inspect relevant access activity rather than leaving security evidence visible only to the service operator.

VSafe separately retains Document Transfer records. Transfer logs can be viewed and downloaded by the account holder, providing a record of transfer activity independently of the temporary document payload.

This supports accountability while allowing the confidential documents themselves to expire from temporary storage.

Security Audit Trail

VervIAM records authentication and protected-access activity required to establish an application-level audit trail.

Security logging is designed to retain sufficient operational context for authentication, token validation, protected endpoint access and abnormal behaviour to be understood during monitoring and investigation.

Human readability is an explicit operational objective: security evidence should be capable of rapid interpretation when investigation is required rather than requiring reconstruction from otherwise disconnected technical events.

Platform Monitoring and Assurance

Verv supplements its application-level audit records with monitoring across the AWS Verv environment.

Custom CloudWatch dashboards provide visibility of authentication behaviour, token-validation outcomes, protected endpoint activity, API errors, request and integration latency and other operational security conditions.

Relevant alerts provide detective controls for conditions requiring investigation, while AWS platform security and monitoring services provide additional infrastructure-level evidence.

The monitoring architecture can be customised and can support integration with an enterprise's existing logging and security-monitoring environment where required.

External Threat Intelligence and Risk Event Detection Alerts (REDA)

Verv operates a daily REDA security-assurance process that consolidates current external threat intelligence with security evidence generated by the Verv environment.

REDA combines reference-sourced external threat and vulnerability intelligence with security evidence generated by the Verv environment. Verv network, WAF, API and authentication evidence is supplied as context to an OpenAI-assisted analytical review, allowing observed activity against Verv services to be assessed against the current external threat environment.

REDA also identifies conditions requiring investigation, including attempts to use protected APIs where the corresponding required VervIAM authenticated-access context cannot be established.

As at the date of this Security Posture assessment, daily REDA monitoring had identified no successful unauthenticated access to a Verv-protected endpoint.

Email Assurance

VervIAM service email is delivered through Amazon SES and uses configured domain email-authentication controls.

DKIM and DMARC provide authentication and reporting for the VervIAM service domain. DMARC reports are delivered automatically and reviewed daily as part of Verv security monitoring.

This provides a continuing evidence stream for the authentication of legitimate VervIAM service email and visibility of authentication anomalies affecting the service domain.

Assurance Through Evidence

Verv does not treat the presence of preventive security controls as sufficient evidence that those controls are operating as intended.

Authentication records, access logs, transfer records, platform telemetry, configured alerts, email-authentication reporting, external threat intelligence and REDA analysis provide complementary evidence from different parts of the security architecture.

This allows security operation to be examined from multiple perspectives and provides evidence for investigation, operational review and continuing improvement.

Enterprise Incident Response Integration

Identity, access control, encryption, data protection and monitoring are fundamental enterprise security controls. VervIAM was architected to apply them across identity, protected application access and data.

Enterprise VervIAM can contribute security evidence to an organisation's wider incident detection and response capability. Authentication and protected-access records, endpoint activity, application and API telemetry, security alerts and other Verv monitoring information can provide contextual evidence for investigation and response.

The logging and monitoring architecture is designed for integration with customer-specific security monitoring, allowing Verv security evidence to contribute to an enterprise's existing incident-response processes rather than operating as an isolated application audit trail.

This becomes increasingly relevant where enterprise transactions include automated or AI-mediated processing. Identity, access control, encryption, data protection and monitoring remain necessary controls regardless of whether an endpoint is operated by a person, conventional application or AI service.

Security Boundary and Claims

Verv security assurance is intended to establish evidence about activity within the systems and security boundaries that Verv controls.

Monitoring cannot establish that compromise is impossible, nor can Verv guarantee the security of a participant's independently controlled device, browser, email service or external enterprise environment.

Security claims are therefore based upon the controls implemented within the Verv architecture and the evidence produced by their operation rather than upon an assertion of absolute security.

Security Assertion

Verv combines privacy-conscious architecture with user-visible auditability and continuing security assurance. Private information is protected and retained according to its operational purpose; account holders can inspect relevant access and transfer records; and application logging, AWS monitoring, configured alerts, email-authentication reporting and daily REDA analysis provide complementary evidence of security activity. Preventive controls are therefore supported by continuing observation and review rather than being assumed to operate effectively without evidence.