

here is the binding of public key and user identity, issued by our web-based authority that utilizes the authentication service offered by firm OA. With which, employees are supposed to show up their SecureID for taking advantage of this toolkit. Notice that, at this stage it is essentially a password-based authentication honored by trusted third-party (firm OA), resulted in a provable identity comes into being for the future uses.

- *Core system (upper right, five blue ovals):* this accomplishes TAAS major functionalities, including identity authentication, privilege validation, credential release, access authorization, and resource management. We discuss these issues in details subsequently. By now, we primarily focus on its constructions. The most striking one is authentication center and capability certificate issuer (*taas_d*), responsible for verifying user identity and its claimed role so as to authorizing a role attribute-based certificate (ticket interchangeably) if succeed. Therefore, user possessing such a certificate is analogous to hold a *proof-carrying* exhibit in hand, with which to request accesses to targets. The second component is a security coordinator (*taas_sec*) for shared key distribution, also in charge of distributed access toward information-flow control, an ultimate end-to-end security property. The third one is a collection of DB oriented API services (*taas_biz*) that mediate the access to sensible data. Others like web front-end (*taas_web*) and databases (*taas_db*) are supplied as usual, particularly for supporting resource management. Here, we omit to present an important component in Figure 1, an authorization library which is the implementation of authorizing framework.

Considering interactions between TAAS and other components in our cloud infrastructure, this is exemplified in lower part of Figure 1 (different colored ovals represent different systems). For instance, imaging a distributed database system (XCUBE in green) launched as a map-reduce job, a client application, to be submitted to TORCA master (in grey). In turn, this job is dispatched as many tasks running on the workers (ES in grey), one of which is taken as XCUBE Master (MS), others as XCUBE workers (TS). They probably have to access XFS system (in pink) for pulling or storing data from which. For security concerns, access control is inevitable while these sorts of accessing requests occur across different component domains in the work-flow. This leads to the interactions with TAAS to perform authentications, represented by some dotted lines in figure.

III. TAAS AUTHENTICATION

A. Authentication Protocol

We design and implement such a novel protocol that combines the merits of PKI-based and Kerberos-based authentication protocols, say offline credential generation and single-sign on (SSO) properties, while harnessing system TCB for trusted delegation. These techniques reasonably mitigate the authentication workload, resulting in a lightweight authentication architecture. We intuitively outline this

protocol by Figure 2 and the key points there, followed by its formalization in Figure 3.

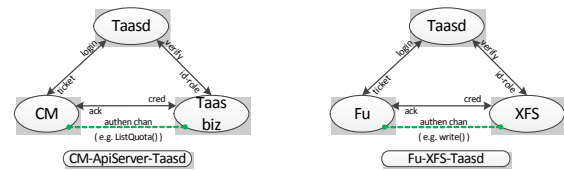


Figure 2. TAAS authentication core procedure.

The left and right diagrams in Figure 2 represent generic client-server authentication scenarios by using TAAS service. Say in the left one, (TORCA) CM intends to authenticate itself to Taas biz service in hope of establishing channel. Recall that our toolkit has already generated necessary credentials offline for client CM to prepare future authentications. This is the start point where CM takes *login* action. In this step, CM's credential is sent to Taasd for validating its identity and claimed privilege assumed by role. Were it succeeded, a ticket and session key will be sent back to CM, which is a PKI-based authentication scheme rolling up. Having received crypto, CM makes its authenticator then concatenates with ticket, assembled together as a new credential standing for its *proof-carrying* identity and capability. Thereafter, credential is delivered to Taas biz for verification, in turn, to be forwarded to Taasd that is the authority appropriate for this job. Later on, verification result echoes back to Taas biz. Based on the answer, it knows whether that requesting channel should be established or not. Observe that, a green dotted line in figure exactly indicates the success of authentication and the buildup of authenticated channel. To this end, it is not surprising to think of the main part of our protocol as Kerberos-based. The true story is that we did more than that, reflected in our proposed novel adaption and optimization for this protocol design, such as, consecutive authentication across domains, group key issues, information flow control and so forth, part of which are addressed below.

In contrast to intuitive illustration above, we also give a formal semantics of this authentication protocol in Figure 3, to help people understand its core procedure more precisely.

B. Some Advanced Topics in TAAS Authentication

- *Access control security vs. information-flow security:* the former enforces authorization merely at access point, however the latter pursues an end-to-end security goal as information propagated over the system. For instance, due to work-flow, user's credential must be exposed at different places where access control is required and conducted. Guaranteeing such data confidentiality and integrity is a notorious hard problem both in theory and practice, to which we develop a security coordinator (*taas_sec*), dynamically regulating credential usage in work-flow whenever extruding its original domain.
- *Proof-carrying authentication and authorization (PCA):* this is a promising technique particularly appropriate for distributed and decentralized access control. A

typical scenario can be found in our work, where an anonymous user may access to XFS data node to execute read-write actions. Due to performance considerations in cloud computing, it is unrealistic for such data nodes to carry out regular access control, instead, a feather-weight verification imposed on user's certificate is expected for simply deriving its capability.

- *Attack models*: as applied in a private enterprise cloud, certain trustworthy is assumed so that we rather than too care about malicious behavior happened in network transition, but focus on replay attack instead, because in this context it is reported to be the most common attacking pattern. We resolve this problem inspired by the idea in Kerberos. Novelty relative to this part shows that sort of address-based attack is prevented without interfering application work-flow, thanks to our security coordinator that safely alters IP in credential.
- *Certificate revocation*: this is a serious problem in pure PKI-based authentication system. In our work, right after early PKI related login procedure, protocol steps into Kerberos-alike processing without such hassles intrinsically. In fact, certificate in TAAS is issued by centralized authority (*taas_d*) that stamps cert with short session time, effectively avoiding related problem.

```

client → taas_d: cred1, where cred1 ≜ certc ⊕ auth1
certc ≜ {idc, Kc+} Kca+
auth1 ≜ {idc, role, Tv, times} Kc-
Tv ≜ time stamp
times ≜ Tb, Te, Tr

taas_d: isOk = let m1 = {cred1|certc} Kca+, m2 = {cred1|auth1} m1|key
in m1|id ≜ m2|id && m2|id ∈ m2|role
keypack ≜ new Kca in {Kss, t} m1|key
ticket ≜ let m = {idc, role, times, IPc, Kss} in {H(m)|Ktaas} K2

taas_d → client: if isOk then return (keypack ⊕ ticket) else return err
client: auth2, where auth2 ≜ let m = {keypack} Kc- in {idc, t'} m1|key

client → server: cred2, where cred2 ≜ ticket ⊕ auth2
server → taas_d: cred3

taas_d: isAuth = let m0 = {cred2|idc} K2, m1 = m0|msg, m2 = {cred2|auth2} m1|key
in H(m0) ≜ {m0|msg} Ktaas+ && m1|id ≜ m2|id && m1|ip ≜ ReqIP

taas_d → server: if isAuth then return (idc, role) else return err
    
```

Figure 3. Formalization of TAAS authentication protocol.

IV. TAAS AUTHORIZATION

The core of this part consists of a policy model and an access decision procedure. To the former, we invent R_ABAC as an effort to reduce the complexity of policy administration and realize flexible fine-grained authorization. To the latter, we implement a generic framework aimed at usability and scalability, akin to Java security architecture.

A. R_ABAC Policy Model

We adapt RBAC₀ as the underpinning, combined with ABAC model for rule-based constraint. To do so is primarily arising from the real needs. For example, whether a user can be authorized to write to a specific XFS file depends on not only her role that represents permissions, but role-associated storage quota that must agree with quota threshold for write enable. Such threshold can be viewed as a role attribute, the

constraint that prescribes minimum quota value for certain operations allowed. Thus, a need for R_ABAC emerged. One instance in practice is presented in Figure 4.

RBAC Policy			Role-attribute Constraint Policy				
Role	Cluster	Permissions	Role	Cluster	Quota	#Files	#Dirs
R ₁	C ₅	(AllAct, Files _{R1})	R ₁	C ₅	20G	3000	200
R ₂	C ₈	(AllAct, Files _{R2})	R ₂	C ₈	40G	6000	400
...

Figure 4. R_ABAC policy example.

B. Implementations

The lower half of Figure 5 delineates main constructs of authorization lib, important interfaces as well as a UML-like time chart to present interactions. As we have seen, these are implemented based on relative standard framework, following PAP, PEP, PDP, et al design scheme, if abstract.

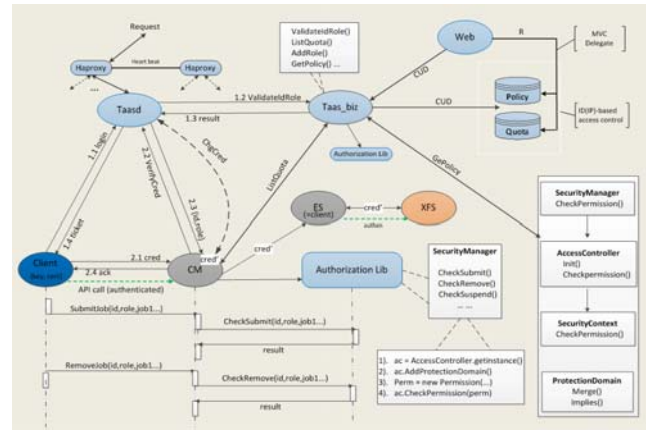


Figure 5. TAAS authorization framework.

V. TAAS RESOURCE MANAGEMENT

In our cloud platform, TAAS is in charging of the resource management not only for identity, role, and policy relativized to security, but for computation and storage quota that as aforesaid, are allocated on the basis of *role*. The central problems existing in such schema are that, one is about how to classify these resources and organize them in a logical hierarchy based on their affiliation; the other is about how to draft application workflows that influence the way of applying resource, processing approval, and auditing results. Obviously, the former concentrates on structuring a well-organized extensible architecture, while the latter emphasizes the flexible and accountable dynamic utilization of resources. Toward the specifications and goals above, we design and implement a management model in Figure 6 below.

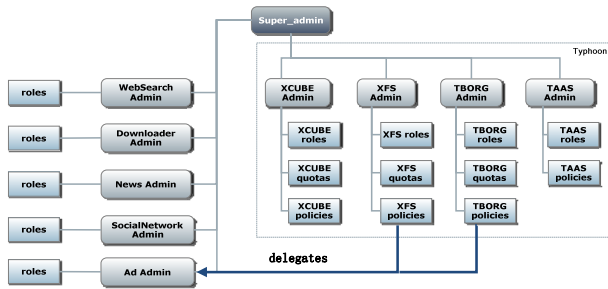


Figure 6. TAAS resource management model.

As seen, the system is logically layered as super-admin, cluster-admin and project-admin in hierarchy, each of which is responsible for managing and maintaining resources pertain to whom, say roles, as well as dynamic resource applying and partial approval at disposal. A worth mentioning point is that we introduce a *delegates* relation from cluster-admin to certain project-admin, for the designated cluster is solely possessed and used by some project. Doing so is to consider resource usability and management flexibility

VI. EXPERIMENTAL RESULTS AND SYSTEM DEPLOYMENT

Since TAAS authentication mechanisms, *login* and *verify*, are essentially implemented in network application layer toward secure communication, we are more concerned with the overhead for creating such an authenticated channel, which reflects in overall response time and throughput comparing with that of non-secure channel. Figure 7 shows the related testing outcome. An exciting statistical result manifests that such overhead is less than 3.1%.

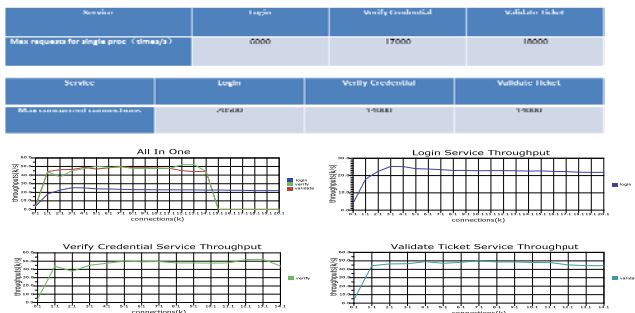


Figure 7. Experiment results of response time and throughput.

To achieve system reliability and localized service, Taas is deployed based on IDC described in Figure 8, where totally four suits of identical services are set up in cities: SZ, XA, TJ, and SH respectively. In case one site crashes, any of other three will take effects to continue serving. In fact, to meet reliability ends, one of our measures beside such redundancy is of an advanced fault tolerance mechanism, in support of regular business processing even if all TAAS services break down.

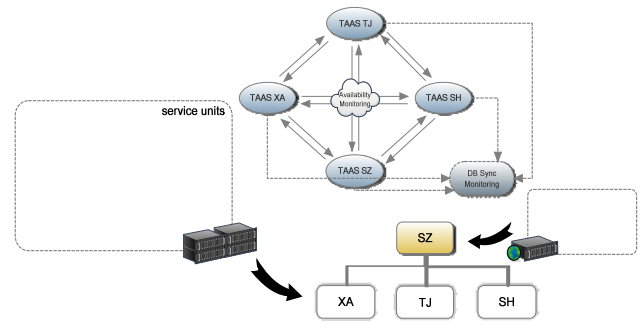


Figure 8. TAAS system deployment and monitoring.

Another diagram showed in Figure 8 depicts monitoring of TAAS service availability and data consistency, which is used to inform operation team with service running status, and notify relative faults in a real-time manner, if happened.

ACKNOWLEDGMENT

We are very grateful to our colleagues and leaders, for their insightful suggestions, continuous help and contributions to this work. They are Guang Yang, Allen Liu, Tao Wu, Bo Li, Hanmei Luo, Dafu Deng, Xin Yao, Jun Chen and Dr. Huican Zhu.

REFERENCES

- [1] B. Lampson, M. Abadi, M. Burrows, and E. Wobber, "Authentication in distributed systems," *ACM Trans. Computer Systems*, 1992.
- [2] A. Appel, E. Felten, "Proof-Carrying Authentication," ACM Conference on Computer and Communications Security, 1999.
- [3] R. S. Sandhu, E.J. Coyne, H.L. Feinstein, C.E. Youman, "Role-Based Access Control Models," *IEEE Computer* 29(2): 38-47, IEEE Press.
- [4] E. Yuan, J. Tong, "Attributed Based Access Control (ABAC) for Web Services," International Conference on Web Services (ICWS'05).
- [5] Dominic Duggan, Ye Wu. "Secure Nested Transactions," IEEE Symposium on Security and Privacy, 2011
- [6] Dominic Duggan, Ye Wu. "Causality and Accountability", Formal Aspects of Security and Trust, LNCS, 2009.
- [7] B. Lampson, "Practical Principles for Computer Security," Software System Reliability and Security, 2006.