Automated Distribution of Out-of-Band Key Material in Virtual Private Networks

David Schatz[1], Hedwig Koerfgen[2] and Guenter Schaefer[1]
[1]*Technische Universität Ilmenau, Germany*
[2]*Universität der Bundeswehr München, Germany*
{*david.schatz, guenter.schaefer*}*@tu-ilmenau.de, hedwig.koerfgen@unibw.de*

Keywords: Quantum-Resistance, Key Exchange, Virtual Private Networks, Business Trips

Abstract: Business Trip Key Exchange (BTKE) is a novel approach to distribute symmetric key material in virtual private networks (VPNs). The main idea is to distribute key material in an automated way during business trips of employees. In this article, we present a protocol for the secure transfer of key material between VPN gateways and mobile devices carried by employees during trips, including aspects like secure persistent storage, and replacement of old key material if required. Using a stochastic travel model representing realistic scenarios, we show that BTKE is able to quickly distribute key material within the VPN. In result, and in combination with multipath key reinforcement (MKR) and post-quantum cryptography (PQC), BTKE is able to provide an additional security layer to realize quantum-resistant VPNs.

1 INTRODUCTION AND MOTIVATION

The ongoing research in quantum computing threatens the security of virtual private networks (VPNs). That is because currently widely deployed schemes to realize *asymmetric cryptography*, such as RSA, ElGamal, and the Diffie-Hellman key exchange, can be broken by sufficiently large quantum computers (Shor, 1997; Proos and Zalka, 2003). Furthermore, alternative solutions for a quantum-resistant key exchange have to be deployed as soon as possible due to the "store now, decrypt later" issue.

One mandatory step is the transition to quantum-resistant asymmetric cryptography, known as post-quantum cryptography (PQC). But while there are promising candidates for a quantum-resistant key exchange, known as key encapsulation mechanisms (KEMs), there is a caveat: Their cryptanalysis, especially regarding implementation security (e.g., side-channel attacks), is not as mature as for traditional schemes. For this reason, and for reasons of providing crypto agility in general, various security agencies around the world currently only recommend to use PQC *in combination* with traditional schemes (French Cybersecurity Agency et al., 2024). However, this approach does not fully solve the "store now, decrypt later" problematic in our opinion: In the long run, it might still turn out that a currently deployed KEM scheme has some inherent weaknesses. Consequently, VPN traffic that has been fully recorded today, might be decrypted in the future if the used KEM scheme turns out to be broken in the future.

1.1 Related Work

In addition to PQC, various other alternatives to distribute symmetric key material have been discussed:

- **Quantum Key Distribution (QKD)**: The security of QKD is based on the laws of quantum physics, especially the *no-cloning theorem* (Wootters and Zurek, 1982). Basically, this allows to *detect* whether key material has been eavesdropped. Unfortunately, there are various practical aspects that currently hinder the deployment of QKD, such as the dependence on a "classically" authenticated channel, a limited reach ($\approx 100\,\text{km}$), high costs, and missing security proofs of implementations (French Cybersecurity Agency et al., 2024).
- **Pre-Shared Keys (PSKs)**: The manual distribution of PSKs (e.g., using trained human key carriers) is usually assumed to be very secure. Unfortunately, it is also very time-consuming and expensive, especially if key material shall be renewed on a regular basis.
- **Multipath Key Reinforcement (MKR)**: The

Preprint; accepted and in press at ICISSP 2026, https://icissp.scitevents.org/

main idea of MKR is to distribute key material *in-band* via different paths (i.e., using VPN connections in our scenario). Then, the different materials are combined at the endpoints do derive the final key (Deng and Han, 2008; Rass and König, 2011; Schatz et al., 2023). As long as not *all* paths are compromised, the resulting key is secure (not known to attackers). However, one limitation of MKR is that it can only provide additional security for a connection between endpoints u and v if there already is an uncompromised path between u and v.

1.2 Contributions of this Article

While MKR is able to efficiently increase the effort for attackers, an initially secure key exchange between some endpoints in a VPN is required. Hence, the main research question of this article is:

> How to ensure that there is an additionally secured (in addition to PQC) *spanning tree* across all endpoints of a VPN, that is, there is one additionally secured path between any pair of endpoints? And how to achieve this in a scalable way, without relying on QKD or manually distributed PSKs?

To answer this question, this article evaluates the idea of a *Business Trip Key Exchange (BTKE)*: Instead of relying on manual distribution of PSKs, BTKE takes advantage of "normal" business trips between sites of an organization (company/agency) to distribute key material in a fully automated way. In this article, we focus on VPNs that are realized by using the IPsec protocol family, but the approach is equally applicable to other networks that consist of both mobile clients and stationary gateways. In this scenario, mobile *VPN clients* (mobile devices of employees) act as *couriers* that automatically receive key material, store it, and distribute it to *VPN gateways* at other sites of an organization. The expected benefit of this approach is mainly motivated by the fact that if enough key materials are distributed this way, even very powerful attackers will have a hard time to intercept all of them. This is especially true because the required communication between devices is designed to take place in *private networks* if possible.

In result, BTKE represents one possible source of key material for recent proposals to use additional key material in the context of IKEv2: The approach suggested in (Schatz et al., 2023), encrypting all IKEv2 traffic using a separate component, is especially suited for BTKE: In contrast to using PSKs within IKEv2, as suggested in (Fluhrer et al., 2020) and (Smyslov, 2025), it protects against implementation flaws in the IKEv2 daemon, and includes a protocol to synchronize the combination of multiple key materials (from BTKE and other sources) over time at the endpoints.

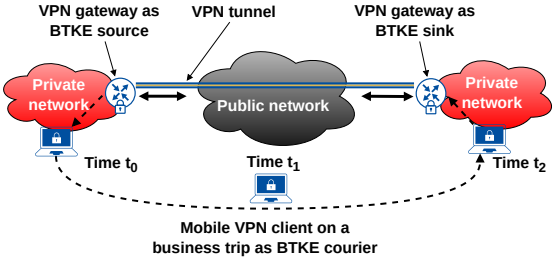

Figure 1: The involved devices and roles of a BTKE.

In the remaining article, we formalize the BTKE and evaluate its practicability and effectiveness. Based on the requirement analysis in section 2, we derive a suitable design in section 3. Especially, this includes a protocol to transfer key material between VPN gateways and mobile devices, and an algorithm to decide on the number of transferred key materials. To evaluate the BTKE in realistic scenarios, we derive a suitable travel model in section 4. Using this model, we perform a quantitative evaluation of the BTKE in section 5. Based on our results, we discuss how BTKE compares to alternative approaches in section 6. We conclude our article in section 7.

2 REQUIREMENT ANALYSIS

We first present a general overview of the BTKE scenario and the underlying threat model. Subsequently, we derive specific requirements for the BTKE.

2.1 Scenario

In total, three roles are involved in a BTKE, as depicted in Figure 1: the *source*, a *courier*, and the *sink*. Source and sink are placed on *VPN gateways*, which secure the communication between private networks at different sites. Couriers are *mobile* VPN clients, which move between different private networks and distribute key material between gateways.

2.2 Threat Model

Our threat model is based on the Dolev-Yao model (Dolev and Yao, 1983): Attackers can compromise some VPN gateways and VPN clients. Furthermore, they can control the network traffic in some of the private networks, that is, eavesdrop, inject, replay, drop, and manipulate packets between gateways and couriers. However, we assume that attackers are

not able to compromise all couriers at all times, because otherwise, BTKE would not be able to provide additional security. Further, we assume that attackers can not compromise all VPN gateways or private networks at all times. In this case, a VPN would fundamentally not be able to protect any user traffic.

2.3 Requirements for the BTKE

We identify the following requirements in the described scenario:

2.3.1 Functional Requirements

The general objective of the BTKE is as follows: If a courier is first located within the private network of a VPN gateway u, and subsequently in the private network of a gateway v, key material shall be transferred from u (via the courier) to v in an automated way. The specific functional requirements are categorized by role:

F1. Upon request by a courier, **a BTKE source** has to generate fresh *key material* and transmit a copy to the courier. Further, it has to persistently store a copy of the key material itself until the other copy has arrived at some sink (BTKE successful). However, it might be necessary to replace the key material on the source in the meantime due to limited storage (BTKE failed).

F2. **A BTKE courier** must be able to detect whether he is located within the private network of a VPN and identify the corresponding gateway. Subsequently, he should deliver stored key material from other private networks to all available BTKE sinks, and in turn request fresh key material from all available sources.

F3. **A BTKE sink** has to receive key material from couriers and use it to additionally secure the VPN connection to the source of the key material.

2.3.2 Security Requirements

Handling of key material has the following security requirements:

S1. **Secure Generation**: The BTKE source must generate fresh key material using a cryptographically secure random number generator.

S2. **Authenticated Transfer**: Optionally, a courier should establish a mutually authenticated channel to each source and sink. The main motivation for this is *availability* (of storage space), not confidentiality: If a private network is compromised, attackers are able to break confidentiality for all actual traffic from/to this private network anyway. Furthermore, even if some individual BTKE exchanges can be observed by attackers in some private networks, the main security benefit of BTKE arises from the assumption that attackers will not be able to observe all exchanges at all times in all private networks. Still, successful mutual authentication may prevent fake BTKE requests to deplete storage resources on uncompromised devices.

S3. **Secure Persistent Storage**: Source and courier must persistently store key material because devices might be rebooted before a BTKE is successful. This storage should be secured to mitigate the risk of key material getting compromised at rest.

S4. **Secure Deletion**: After a successful BTKE, key material must be securely deleted from all involved devices to prevent a later compromise.

2.3.3 Non-functional Requirements

We further state the following non-functional requirements for the BTKE.

N1. **Limited Storage**: It should be possible to limit the required storage space on BTKE sources and couriers by employing a suitable replacement strategy for old key material.

N2. **High Coverage**: If a courier moves between private networks of gateways $v_1, \ldots, v_n$, many pairs $(v_i, v_j), i \neq j$ should benefit from a successful BTKE.

N3. **Scalability**: BTKE should scale both with the total number of gateways in the VPN and with the total number of mobile clients.

3 DESIGN OF THE BTKE

In this section, we derive a design for the BTKE. We focus on a suitable protocol to generate, transfer, and replace (if required) key material. However, two other aspects that have to be considered in practice heavily depend on the specific implementation and topology of the VPN:

1. **Discovery of Local VPN Gateways**: To fulfill requirement F2, in the best case, a courier is able to find his local gateway via traditional routing information (e.g., by using Traceroute). More complex topologies (e.g., nested private networks with multiple middleboxes such as firewalls between courier and gateway) might require more sophisticated discovery mechanisms.

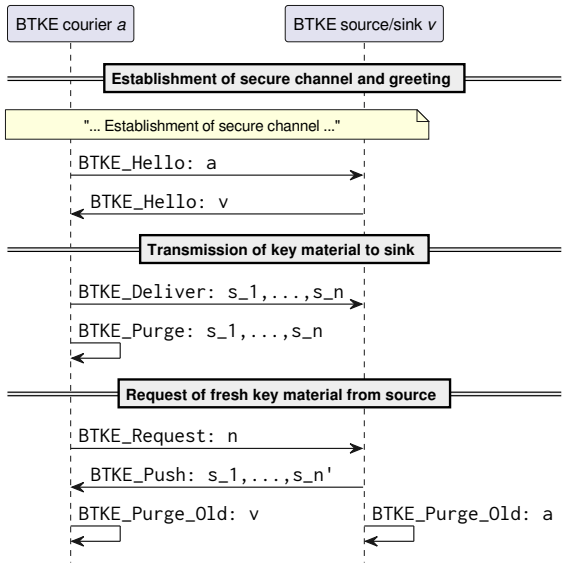

Figure 2: BTKE transmission protocol, with a notation inspired by Remote Procedure Calls (Operation: Parameter).

2. **Secure Communication between Courier and Source/Sink**: If the courier can act as a VPN endpoint himself (e.g., to connect to the VPN from outside any site), the communication within the private network may also be secured by utilizing the IPsec protocol family. Otherwise, suitable alternatives (e.g., TLS) might have to be used. As discussed in subsubsection 2.3.2, the overall security benefit of BTKE for the VPN *does not rely* on this secure channel. Still, it may act as a "security in depth" measure to further increase attack costs.
3. **Secure Persistent Storage**: In some cases, VPN gateways and mobile clients within an organization might already use a mandatory disk encryption. In other cases, at least the BTKE key material should be protected by means of disk encryption on gateways and couriers (requirement S3).

Consequently, these aspects are not further discussed here. In the following subsections, we present the specific protocol to transfer key material, and an algorithm to replace old key material if required. For this, we assume that every VPN gateway and every courier has a *unique ID*. We further expect this ID to be checked during mutual authentication between gateways and couriers (requirement S2).

3.1 Transmission Protocol

As depicted in Figure 2, the transmission protocol between a courier a and a gateway v is divided into three phases. During the protocol, the gateway first acts as sink, and subsequently as source:

1. First, the secure communication channel is established, where the parties mutually authenticate each other using their identities a and v.
2. Next, the gateway acts as sink. For each of the n sources from which a courier stores key materials, exactly one key material is sent to the sink (requirements F2, F3).
3. In the last phase, the gateway acts as source and the courier requests n fresh key materials (requirement F2). Depending on its current storage utilization, the gateway securely generates and transmits $n' \leq n$ key materials (requirements F1, S1). The order of the individual key materials in this *key set* is important for the replacement strategy discussed in the next subsection. Finally, old key materials which are no longer required are securely deleted by both parties (requirement S4).

3.2 Replacement Strategy

Due to limited storage space (requirement N1) it might be required to replace old, unused BTKE key material. To design a suitable replacement strategy, we assume that some (configurable) properties of the VPN are known:

- A globally known upper bound for the total number $n_{\max}$ of VPN gateways.
- The available storage capacity $c(v)$ (in number of key materials) for each gateway v, and analogously, the available storage capacity $c(a)$ for each courier a (locally known).

The essential questions during the third phase of the transmission protocol (see Figure 2) are:

1. How many key materials (number n) should the courier a request?
2. How many key materials (number $n' \leq n$) should the source v actually generate? And which key materials should it replace if required?
3. Which key materials should a delete if required?

Our proposed replacement strategy aims to provide a fair allocation of the available storage space on a gateway to all known couriers, and vice versa. Furthermore, if both a source and a courier have to delete key material from the same key set, they should delete the same key materials (instead of different ones, which would unnecessarily reduce the number of key materials known to both, see requirement N2). Last but not least, it is preferable to replace older key material compared to newer key material.

One easy implementation to meet these objectives is to use a hierarchical storage layout as depicted in

Preprint; accepted and in press at ICISSP 2026, https://icissp.scitevents.org/

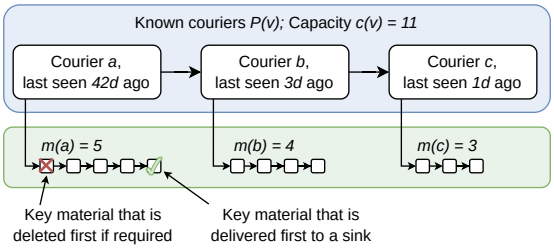

Figure 3: Example for the layout to persistently store key material on BTKE source v (analogously on couriers). In the example, v knows $|P(v)| = 3$ couriers, storing $5, 4$, and 3 key materials, respectively. Consequently, its total capacity of $c(v) = 11$ is exceeded by $d = 1$ and the marked key material shared with a is deleted.

Figure 3. The upper layer stores the known communication *peers* (i.e., couriers in case of gateways and vice versa), primarily sorted descending by the number of shared key material and secondarily sorted descending by the time since the peer was last seen. The lower layer further describes a deterministic ordering of key materials within key sets. This dictates the order in which couriers deliver key material to sinks (back to front) and order of deletion (front to back).

The actual algorithm based on this layout works as follows: First, the courier a determines the number n of key materials as follows, with $P(a)$ being the set of all BTKE sources he knows:

$$n = \min\left\{\left\lceil {}^{c(a)}/_{|P(a)|}\right\rceil, n_{\max} - 1\right\} \tag{1}$$

This aims to provide a fair allocation of storage space to all known sources and $n_{\max} - 1$ is an upper bound for the number of possible sinks for the requested key material. Analogously, the source v determines how many key materials to actually generate, with $P(v)$ being the set of all couriers it knows:

$$n' = \min\left\{\left\lceil {}^{c(v)}/_{|P(v)|}\right\rceil, n\right\} \tag{2}$$

Note that this does not generate more key materials than originally requested by the courier (n).

Next, let $m(b)$ be the number of stored key materials for each known courier $b \in P(v)$, with $m(a) = n'$. Consequently, the number d of key materials that have to be deleted on the source is:

$$d = \left(\sum_{b \in P(v)} m(b)\right) - c(v) \tag{3}$$

If $d > 0$, the source repeats the following procedure d times: It selects the courier which is in the front according to the sorting of the hierarchical storage layout and deletes exactly one key material of the corresponding key set. Then, the list is sorted again before moving on to the next iteration. During this procedure, the currently requesting courier a is skipped, because the shared key set with a has been completely replaced (with a suitable new size n') anyway. Analogously, the courier frees up storage after receiving the n' fresh key materials.

Note that the replacement strategy is dynamic (apart from a rough upper bound for the number of gateways), and consequently scales well with the number of gateways/couriers (requirement N3).

4 MODELING BUSINESS TRIPS

For the quantitative evaluation, we assume that BTKE is used for exactly one VPN of one organization with multiple sites. To simulate the distribution of key material over time, we require movement profiles for the employees of the (simulated) organization. To obtain such profiles, we considered two approaches:

1. **Empirical Movement Profiles**: We first searched for a dataset containing empirical movement profiles from real organizations. Unfortunately, our search using the databases OECD Data Explorer (OECD, 2025), SocioPatterns (ISI Foundation et al., 2025), and GESIS Search (GESIS, 2025) did not find a suitable dataset. Furthermore, studies such as "National Travel Survey 2024" (U.K. Department for Transport, 2024), "National Household Travel Survey" (U.S. Department of Transportation, 2022), "Mobilität in Deutschland" (Follmer, 2025) focus on the means of transport, not on detailed movement profiles within organizations.
2. **Stochastic Models**: Existing models for artificially generating movement profiles do not explicitly address business trip (e.g., (Chang et al., 2014; Shahzamal et al., 2014)), or only model business trips in an aggregated form (e.g., on the granularity of countries (Kulendran and Wilson, 2000)).

Consequently, we assume the following parameterizable stochastic model to generate movement profiles. By adjusting the parameters of the model, different organizational structures may be simulated. The objective is to realistically model the following aspects of business trips:

1. **Location of Sites**: The geographical location of sites is required to model the *distance* between them, which naturally should influence the likelihood of business trips between sites.
2. **Number of Employees at Sites**: The number of employees at each site also influences the likelihood of business trips between sites.
3. **Business Trips**: Trips between sites are modeled individually for each user, and require sub-models for frequency, destination, and duration.

4.1 Zipf's Law as Model Foundation

Zipf's law describes a typical distribution of real world (and especially human) processes: Some few events occur very often, whereas a majority of events occur rarely. This phenomenon is also related to the *Pareto principle* (also known as the "80/20 rule") (Newman, 2005). Originally, it was used to describe the relative frequency of words in (longer) texts (Zipf, 1949). However, it is now also used in sociological and economical models, for example, to model the population of cities within a country (Zipf, 1949; Newman, 2005), or to model human mobility (González et al., 2008). Our model uses Zipf distributions to model the size of sites (analogously to city population) and to model the destination of business trips (analogously to general human mobility).

The *Zipf distribution* is a discrete probability distribution over a set A with $|A| = n$ elements, where a unique *rank* $r(a), 1 \leq r(a) \leq n$ is assigned to each element $a \in A$. The definition of the distribution is as follows (using the random variable X):

$$Pr(X = a) = \frac{1}{H_{n,s}} \frac{1}{r(a)^s}, \text{ with } H_{n,s} = \sum_{i=1}^{n} \frac{1}{i^s} \quad (4)$$

Here, $H_{n,s}$ normalizes the sum of probabilities to 1 (generalized harmonic number). The parameter s of the Zipf distribution controls its heterogeneity. While $s = 0$ degrades the distribution to a uniform distribution, $s = 1$ reflects Zipf's law in its original form. Larger values for s result in elements with a high rank to only have a negligible selection probability.

4.2 Modeling Sites

The location of n sites of an organization are modeled based on the geographical location of real cities (in the form of an input dataset). For this, we use a uniform distribution to assign n sites to the cities of the underlying dataset. Consequently, the cities in the input dataset should reflect the geographical extent of an organization that shall be simulated.

Subsequently, m employees are assigned to the n sites (as their main worksite) by using a Zipf distribution. The rank of each site is determined by the population of the corresponding city. That is, on average, there will be some large sites (located in cities with a large population), and many smaller sites. Other parameters during this process are a minimum number m_{min} of employees at each site, and the parameter s_{e} for the underlying Zipf distribution.

4.3 Modeling Business Trips

Individual business trips are modeled as described in the following subsections.

4.3.1 Frequency of Business Trips

The frequency of business trips of employees is modeled via independent *Poisson processes*. That is, business trips of the same employee are also independent, and the behavior does not change over time. For simplicity, we only model workdays.

The total frequency λ of business trips across the organization (average new business trips per day) is allocated to all employees using a Zipf distribution. Consequently, the frequency $\lambda(a)$ for business trips of an individual employee a is calculated as follows, with s_{t} being the parameter of the Zipf distribution, and all valid ranks $1 \leq r(a) \leq m$ assigned randomly:

$$\lambda(a) = \frac{1}{H_{m,s_{\text{t}}}} \frac{1}{r(a)^{s_{\text{t}}}} \lambda \quad (5)$$

4.3.2 Modeling Destinations

The probability to select a specific site as business trip destination should depend on both the distance to the destination and the size of the destination. While some gravity models, such as (Zipf, 1946; Cabanas-Tirapu et al., 2025), directly weight these two aspects in a numerical way, we use an approach that avoids working with absolute values here, because, for example, it is hard to compare "1000 km" with "100 employees". Instead, for each employee a, all possible destinations s are sorted by increasing distance to the main worksite of a on the one side (rank $r_{a,1}(s)$), and sorted by decreasing size on the other side (rank $r_{a,2}(s)$). The final rank for a selection of the destination according to a Zipf distribution with parameter s_{d} is calculated as:

$$r_a(s) = \alpha r_{a,1}(s) + (1 - \alpha) r_{a,2}(s) \quad (6)$$

Here, $\alpha = 1$ reflects a selection based on distance only, whereas $\alpha = 0$ reflects a selection based on size only. Note that our model does not model "circular" trips: Employees always return to their main worksite before starting the next business trip.

4.3.3 Modeling Duration of Business Trips

We model the duration of business trips based on the distance between the main worksite of an employee and his destination: For distances up to 250 km, an employee returns to its main worksite on the same day. For distances between 250 km and 500 km, the

return journey takes place on the next day. For distances above 500 km, the return journey takes places two days later.

5 QUANTITATIVE EVALUATION

In this section, we evaluate BTKE using a *discrete event simulation* in `omnetpp` (Varga and Hornig, 2008). In accordance with our main research question, we focus on the ability of BTKE to ensure that there is an additionally secured *spanning tree* between sites. To answer this question, we evaluate the impact of the various parameters of our business trip model.

5.1 Simulation Setup

Each employee is modeled as a `cSimpleModule` in `omnetpp`, which triggers business trips according to a Poisson process as discussed in subsubsection 4.3.1. This process is paused during a business trip, which results in a slightly lower total frequency compared to the value of the parameter λ. The following subsections further present the default parametrization of the simulation and the evaluated metrics.

5.1.1 Simulation Parameters

The geographic location of sites are selected uniformly at random from a dataset[1] containing all German cities with a population of more than 1000 people. At each site, exactly one VPN gateway is simulated. The size of key material (including an unique ID) is assumed to be 512 bit. Table 1 summarizes all relevant parameters.

Symbol	Default	Parameter
n	100	Number of sites
m	$50n$	Number of employees
$m_{\min}$	5	Minimum size of sites
s_e	1	Assigning employees to sites (Zipf distribution)
λ	$5\,\mathrm{d}^{-1}$	Total business trip frequency (across all employees)
s_t	1	Distributing trip frequency to employees (Zipf distribution)
α	0.5	Trip destinations (weighting of distance vs. size)
s_d	1	Trip destinations (Zipf distribution)
$n_{\max}$	20000	Upper bound for the number of VPN gateways
k	1 GB	Storage capacity of gateways
l	1 GB	Storage capacity of couriers
	1000 d	Simulated time

Table 1: Simulation parameters.

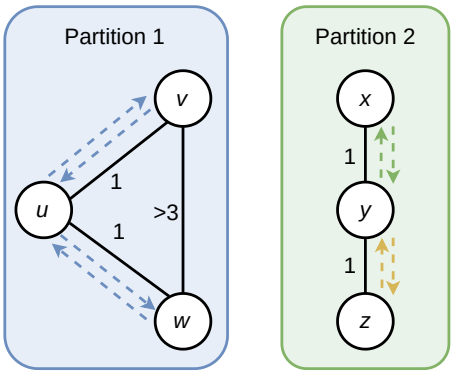

Figure 4: BTKE overlay graph example. Solid lines represent the edges E of the overlay graph $G = (V, E)$, annotated with their age (unitless for the sake of the example). Dashed arrows represent the movement profile of three different employees (three different colors).

5.1.2 Metrics

For evaluation, we model the BTKE as undirected *BTKE overlay graph* $G = (V, E)$: The node set V represents all VPN gateways (and consequently all sites). An edge $\{u, v\} \in \binom{V}{2}$ is added as soon as there has been one successful BTKE between u and v. Further, we assign an $\mathrm{age}(e)$ to each edge $e \in E$. This age represents the age of the key material that has been transferred from source to sink, that is, the time from generation to delivery. If multiple BTKEs have been successful on the same edge e, we update $\mathrm{age}(e)$ to always represent the *youngest* key material that has been transferred. Using this graph representation, we evaluate four metrics:

1. The *number of connected components* (also called *partitions* here) in the graph G over time.
2. The *time to convergence* till the graph G is connected. From this point in time on, a spanning tree exists and MKR may be used to additionally secure *all* pairwise connections within the VPN.
3. The *diameter* of G at the time of convergence, that is, the length of the longest shortest path in G between any two nodes. A smaller diameter implies better starting conditions for MKR.
4. The *average age* across all edges at simulation end, as defined by $\mathrm{age}(e)$, $e \in E$. If couriers are not initially compromised, a shorter storage duration decreases the risk of compromise.

An example for an BTKE overlay graph is shown

[1] https://public.opendatasoft.com/explore/dataset/geonames-all-cities-with-a-population-1000

in Figure 4. In the example, business trips of three different employees generated a total of five edges in the overlay graph G. The key material on four of the edges has been transmitted *directly* during the same business trip, with an assumed travel time of 1 time unit between any pair of nodes. The key material on edge $\{v,w\}$ has been transmitted *indirectly* over the course of two business trips of the same employee, and consequently has a higher age (at least 3 time units for travelling, plus the time between the two business trips). The graph has a total of two partitions, with a diameter of one in the left partition, and a diameter of two in the right partition. If, for example, another business trip between nodes v and x would take place now, the two partitions would be merged (= time of convergence), resulting in a total diameter of four (paths from z to u and from z to w).

5.2 Experiments and Results

Now we present the performed experiments and discuss the results. If not stated differently, all simulation parameters were set to their default value, see Table 1. For all parameters s of the involved Zipf distributions, we also explore some extreme values, ranging from $s = 0$ (degrades to a uniform distribution) to $s = 2$. This is suited to get a general estimation for the effectiveness of the BTKE. We assume that results for real organization would be somewhere in between these extreme cases.

For each evaluated combination of parameter values, we performed 32 independent simulation runs with different seeds for pseudo-random number generation. Result figures show the average value (of the metrics) across all simulation runs, combined with the minimum and maximum value as error bars. Lines connecting the average values are only for visual aid and do not indicate continuous measurements.

5.2.1 Business Trip Frequency

In our first experiment, we studied the impact of the total frequency of business trips on the effectiveness of BTKE. The number of employees is constant for this experiment, that is, increasing λ directly increases the trip frequency for each employee. Our results are shown in Figure 5.

As expected, increasing λ results in a faster time to convergence and younger key materials. Using a low frequency of $\lambda = 1\,\mathrm{d}^{-1}$, convergence was not achieved within the 1000 simulated days in some simulation runs. A positive result of this experiment is that the diameter of the BTKE overlay graph was around 4, independent of the frequency λ (not shown for brevity).

5.2.2 Number of Employees

For a constant travel frequency λ, having more employees implies that each individual employee does travel less frequent. Our results show that this does not have any significant impact on the time to convergence (not shown for brevity). However, there was a significant difference regarding the age of key material, as shown in Figure 6.

When setting $s_t = 1$, having less employees implies that an individual employee a travels more frequently. Consequently, key material that is transferred between two sites over the course of more than one business trip is younger (e.g., a first travels from his main worksite to site u, returns, and later travels to site v to eventually transfer key material from u to v). These indirectly transferred key materials explain why, in this scenario, key material tends to be younger when having less employees, as depicted in Figure 6, left-hand side.

Setting $s_t = 0$ (each employee travels with the same frequency), interestingly shows a different trend, see Figure 6, right-hand side: While each individual employee travels less frequent when increasing m, more key materials are delivered in a single direct trip, which only takes a few hours in our model. Consequently, the average age decreases significantly.

5.2.3 Allocation of Travel Frequency

In this experiment, we further evaluate the impact of the parameter s_t, which controls the allocation of the total traffic frequency to individual employees. Our results are shown in Figure 7.

A more homogeneous allocation of the frequency (lower values for s_t) results in a faster time to convergence. That is because more employees cover more destinations in our model (where destinations are also selected via a Zipf distribution). On the other side, if a majority of trips are mainly performed by a few employees, trips take place between fewer different pairs of sites (fewer main worksites, and fewer possible destinations). But even for $s_t = 2$, convergence is still possible, and the number of partitions in the BTKE overlay graph still decreases rapidly. Furthermore, another practical aspect which is not covered within our model is that even very small sites will usually be involved in business trips on a regular basis (although the absolute numbers may be not as high).

On the flip side, the diameter of the graph at the time of convergence seems to decrease when there is a more heterogeneous allocation. This may actually be explained by the increased time to convergence itself. That is, the diameter within individual partitions has more time to decrease till all partitions are merged

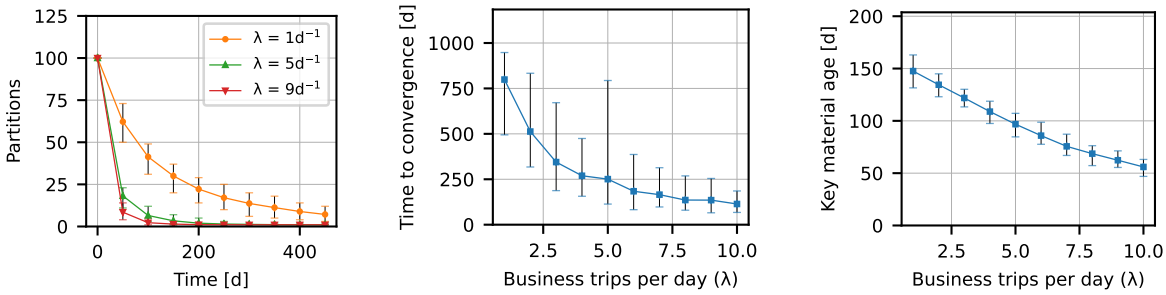

Figure 5: Impact of the total business trip frequency λ.

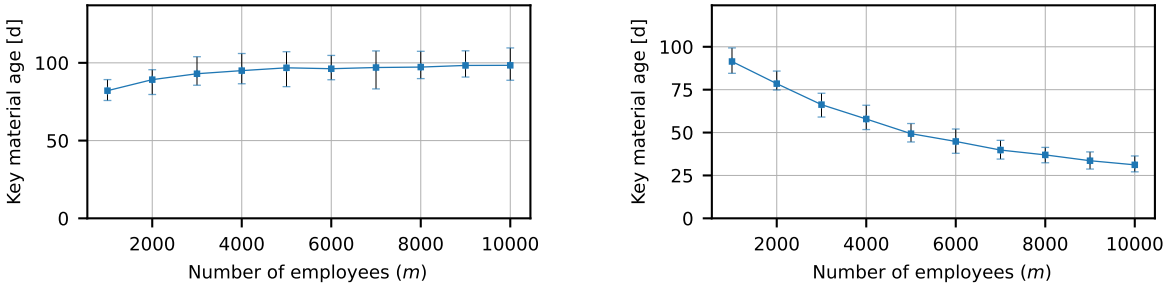

Figure 6: Impact of the number m of employees on the age of key material when distributing the total business trip rate using a Zipf distribution with parameter $s_t = 1$ (left figure), respectively using a uniform distribution, i.e., $s_t = 0$ (right figure).

into one, also resulting in a lower diameter overall.

Our results for key material age do not show a clear trend when varying s_t, not shown for brevity.

5.2.4 Selection of Business Trip Destinations (Zipf Distribution)

Next, we evaluate the impact of the Zipf distribution while selecting the destination of a business trip (parameter s_d). Our results are depicted in Figure 8. As expected from our previous findings, a more homogeneous selection is beneficial for the time to convergence and for the average age of key materials, because more key material is directly delivered during a single trip. The small decrease in key material age when using a very heterogeneous selection ($s_d = 2$) may be explained as follows: First, this scenario also leads to more of the *relevant* key material (i.e., the youngest for each edge in the BTKE overlay graph) being exchanged directly. Second, these directly exchanged keys are also younger in this case because close sites (i.e., faster arrival) are especially likely to be selected. Regarding the diameter of the BTKE overlay graph, we did not observe any significant correlation with the parameter s_d (not shown).

5.2.5 Selection of Business Trip Destinations (Distance vs. Size)

When selecting a trip destination based on the Zipf distribution with parameter s_d, the actual rank for each site is based on both the distance (preferably close), and the size of the destination (preferably many employees). In this experiment, we evaluate the impact of the parameter α, which controls which of these two aspects is weighted stronger (equally for $\alpha = 0.5$).

In Figure 9, it can be seen that preferring close destinations (larger values of α) results in a slightly better time to convergence, and younger key material. This matches our previous results regarding the benefit of more homogeneous behavior, because in our model, the distribution of distances between sites is more homogeneous compared to their sizes.

However, we observe a different trend regarding the diameter at the time of convergence: If the selection process prefers close destinations, there will not be as many "shortcuts" in the BTKE overlay graph between distant sites, which could otherwise decrease the diameter.

5.2.6 Assignment of Employees to Sites

The assignment of employees to sites uses a Zipf distribution. Our results show that the impact of its pa-

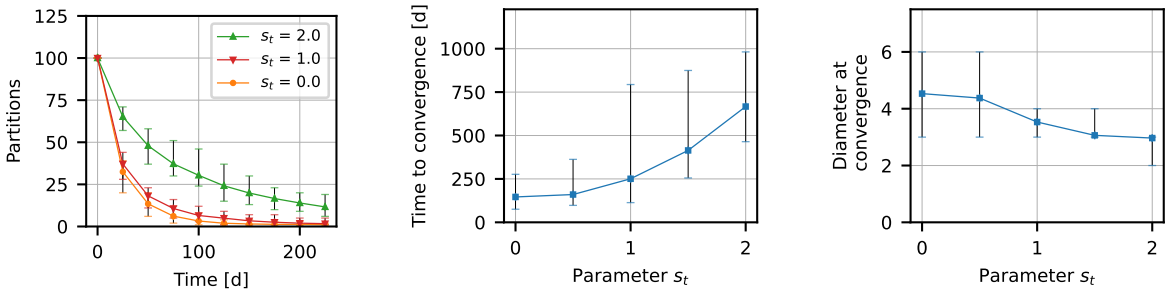

Figure 7: Impact of the parameter s_t.

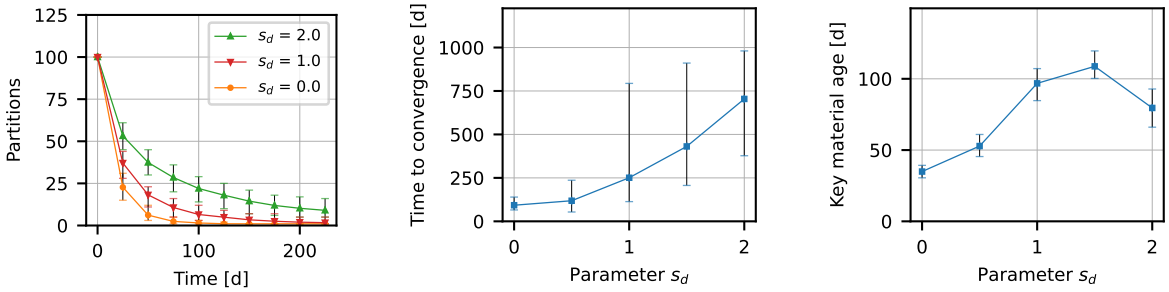

Figure 8: Impact of the parameter s_d.

rameter s_e is minor, not plotted for brevity. And they match our observations from the previous experiment: More homogeneous sizes of sites implicitly result in the distance between sites being more important when selecting sites. Consequently, homogeneity is slightly beneficial for time to convergence and for younger key material, but results in a slightly larger diameter.

5.2.7 Number of Sites

When increasing the number n of sites without increasing the total business trip frequency, the effectiveness of BTKE is significantly reduced (not shown for brevity). However, a more realistic experiment is to also proportionally increase the trip frequency when increasing n. Consequently, we set $\lambda = n/20\,\mathrm{d}$ for this experiment. The results are shown in Figure 10. Unfortunately, both the time to convergence and key material age still degrade with increasing n. Nevertheless, the diameter of the BTKE overlay graph is at a constant low level (≈ 4, not shown for brevity) in all cases. Furthermore, the total number of partitions can still be quickly reduced even for very large n.

5.2.8 Storage Space

Simulating a large decentralized organization with $n = 1000$ sites (and consequently gateways) and $m = 5000$ employees, we were able to reduce the storage capacity down to $k = l = 5\,\mathrm{MB}$ before observing any failed BTKE due to premature replacement of key material. And even for $k = l = 1\,\mathrm{MB}$, the BTKE results in an additionally secured spanning tree in all cases.

5.3 Summary of our Quantitative Evaluation

Our quantitative evaluation showed that BTKE is a robust approach to ensure the existence of an additionally secured spanning tree in almost all studied scenarios, even when using extreme values for the parameters of the business trip model. Varying the total trip frequency and the general homogeneity of travel behavior showed the most impact on the results, with higher frequencies and more homogeneous behavior being the most beneficial for BTKE. In foresight of combining BTKE with MKR, our results show very good conditions for MKR (quick reduction of the number of partitions and a small diameter in the BTKE overlay graph), even when the travel behavior is unfavorable. In practice, a long time to convergence and old key material are only expected to be a problem for very large organizations with a relatively low travel frequency. The remaining parameters of our travel model only slightly impact the effectiveness of BTKE. Last but not least, it is very pleasant to see that BTKE even works under very restricted storage limitations on devices (1 MB).

Preprint; accepted and in press at ICISSP 2026, https://icissp.scitevents.org/

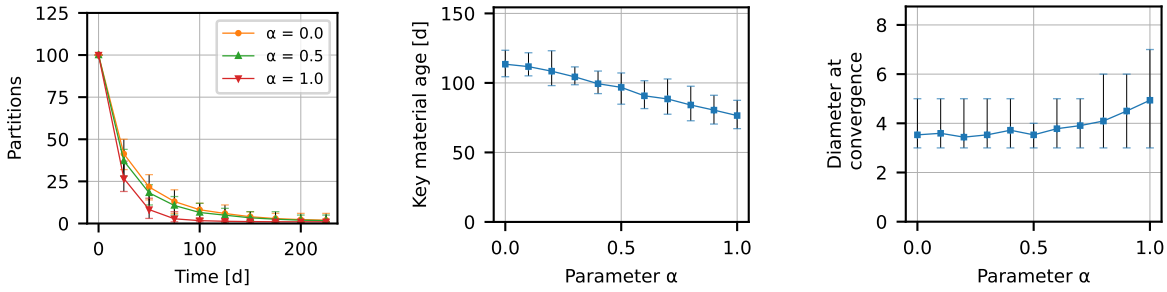

Figure 9: Impact of the parameter α.

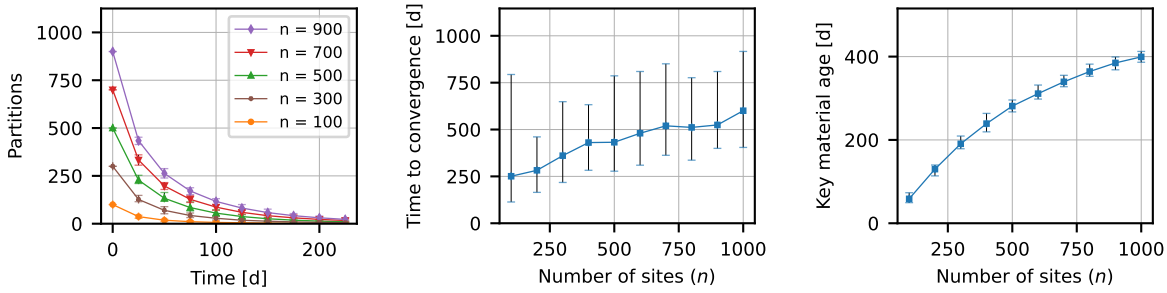

Figure 10: Impact of the number n of sites with a proportional total business trip frequency of $\lambda = n/20\text{d}$.

6 COMPARISON TO OTHER APPROACHES

This section provides a brief qualitative discussion how BTKE compares to other quantum-resistant key exchange methods in the context of VPNs. Note that a transition to PQC is mandatory anyway, so it is not part of this comparison. We solely focus on approaches that may provide *additional* security alongside PQC, with the main candidates being manually distributed PSKs, QKD, and MKR.

Compared to manually distributing PSKs, the recurring costs of BTKE are significantly lower: While BTKE only requires a few MB of disk space on VPN gateways and clients, manually distributing PSKs requires human interaction (transferring key material between VPN gateways and carrier devices, such as smart cards) and causes additional travel costs. The initial costs for deploying BTKE are also low (small software update for existing devices). And while it may be argued that it might be possible to better protect one individual key material during manual transfer (e.g., when using specially trained human key carriers), the overall effort required to intercept *all* key materials exchanged by BTKE is very high due to their sheer quantity. Still, manually distributing some PSKs is a good option to strategically merge partitions in the BTKE overlay graph (which may be integrated into the VPN monitoring) for faster convergence.

The costs of deploying QKD are also significantly higher compared to BTKE. Especially the initial costs for QKD devices are high at the time of writing. Recurring costs might also be high, for example, if dark fibers have to be rented. Furthermore, QKD cannot be deployed for every connection in large VPNs due to current range limitations (French Cybersecurity Agency et al., 2024). And while the security of QKD is very high in theory, an authenticated channel is required beforehand (e.g., by means of manually distributed PSKs), and security in practice may be lower due to side-channel attacks (French Cybersecurity Agency et al., 2024). Still, QKD might in future also provide a way to strategically secure individual VPN connections.

Like BTKE, MKR is a very cost-effective way to additionally secure VPN connections and should also be deployed in VPNs, alongside PQC, in our opinion. However, it requires an additionally secured spanning tree within the VPN to unfold its full potential. As shown by our quantitative evaluation, BTKE can provide such a spanning tree in many realistic scenarios.

7 CONCLUSION

BTKE is a robust, resource efficient, and practical approach to additionally secure VPN connections in a quantum-resistant manner (in addition to a mandatory

transition to PQC). It is able to realize an additionally secured spanning tree within an appropriate time frame in various different scenarios. This is especially true for organizations with many homogeneous sites and many employees that perform business trips on a regular basis, for example, decentralized agencies and companies of national scale. And even in scenarios with a high time to convergence (many sites, few business trips) the resulting diameter in the BTKE overlay graph is small, a good starting point for a combination with MKR. The efficient handling of resources (using existing hardware, low storage requirements) also allows for an easy deployment of BTKE in practice.

As future work, more sophisticated extensions of the business trip model could be studied. For example, "circular" trips, explicit modeling of home offices as main worksites, and individualized travel profiles for employees.

ACKNOWLEDGEMENTS

This research is funded by dtec.bw – Digitalization and Technology Research Center of the Bundeswehr [project MuQuaNet]. dtec.bw is funded by the European Union - NextGenerationEU.

REFERENCES

Cabanas-Tirapu, O., Danús, L., Moro, E., Sales-Pardo, M., and Guimerà, R. (2025). Human mobility is well described by closed-form gravity-like models learned automatically from data. *Nature Communications*, 16(1):1336.

Chang, J. S., Jung, D., Kim, J., and Kang, T. (2014). Comparative analysis of trip generation models: Results using home-based work trips in the Seoul metropolitan area. *Transportation Letters*, 6(2):78–88.

Deng, J. and Han, Y. (2008). Multipath key establishment for wireless sensor networks using just-enough redundancy transmission. *IEEE Transactions on Dependable and Secure Computing*, 5(3):177–190.

Dolev, D. and Yao, A. C. (1983). On the security of public key protocols. *IEEE Transactions on Information Theory*, 29(2):198–208.

Fluhrer, S., Kampanakis, P., McGrew, D., and Smyslov, V. (2020). Mixing preshared keys in the Internet Key Exchange Protocol Version 2 (IKEv2) for post-quantum security. Technical report. https://rfc-editor.org/rfc/rfc8784.

Follmer, R. (2025). Mobilität in Deutschland - MiD Kurzbericht. Technical report.

French Cybersecurity Agency, Federal Office for Information Security, Netherlands National Communications Security Agency, and Swedish National Communications Security Authority, Swedish Armed Forces (2024). Position paper on Quantum Key Distribution. Technical report. https://bsi.bund.de/SharedDocs/Downloads/EN/BSI/Crypto/Quantum_Positionspapier.html.

GESIS (2025). GESIS Search. https://search.gesis.org.

González, M. C., Hidalgo, C. A., and Barabási, A.-L. (2008). Understanding individual human mobility patterns. *Nature*, 453(7196):779–782.

ISI Foundation, CNRS Centre de Physique Théorique, and Bitmanufactory (2025). Datasets obtained through the SocioPatterns sensing platform. http://sociopatterns.org/datasets.

Kulendran, N. and Wilson, K. (2000). Modelling business travel. *Tourism Economics*, 6(1):47–59.

Newman, M. (2005). Power laws, Pareto distributions and Zipf's law. *Contemporary Physics*, 46(5):323–351.

OECD (2025). OECD Data Explorer. https://data-explorer.oecd.org.

Proos, J. and Zalka, C. (2003). Shor's discrete logarithm quantum algorithm for elliptic curves. *Quantum Information and Computation*, 3(4):317–344.

Rass, S. and König, S. (2011). Indirect eavesdropping in quantum networks. In *Proceedings of the 5th ICQNM*, pages 83–88.

Schatz, D., Altheide, F., Koerfgen, H., Rossberg, M., and Schaefer, G. (2023). Virtual private networks in the quantum era: A security in depth approach. In *Proceedings of SECRYPT 2023*, pages 486–494.

Shahzamal, M., Parvez, M. F., U Zaman, M. A., and Hossain, M. D. (2014). Mobility models for delay tolerant network: A survey. *International Journal of Wireless & Mobile Networks*, 6(4):121–134.

Shor, P. W. (1997). Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SICOMP*, 26(5):1484–1509.

Smyslov, V. (2025). Mixing preshared keys in the IKE_INTERMEDIATE and CREATE_CHILD_SA exchanges of the Internet Key Exchange Protocol Version 2 (IKEv2) for post-quantum security. Technical report. https://rfc-editor.org/rfc/rfc9867.

U.K. Department for Transport (2024). National travel survey 2024. Technical report. https://gov.uk/government/statistics/national-travel-survey-2024-technical-report.

U.S. Department of Transportation (2022). National household travel survey. https://www.fhwa.dot.gov/policyinformation/nhts.cfm.

Varga, A. and Hornig, R. (2008). An overview of the OMNeT++ simulation environment. In *Proceedings of ICST 2008*.

Wootters, W. K. and Zurek, W. H. (1982). A single quantum cannot be cloned. *Nature*, 299(5886):802–803.

Zipf, G. K. (1946). The P1 P2/D hypothesis: On the intercity movement of persons. *American Sociological Review*, 11(6):677.

Zipf, G. K. (1949). *Human Behavior and the Principle of Least Effort: An Introduction to Human Ecology*.